Math 8 — Functions and Modeling

Neil Donaldson

Spring 2025

Introduction

This course aims to refresh and reinforce the conceptual foundations behind several topics commonly encountered in grade-school mathematics. The job of a teacher is often one of *selection*: choosing examples and explanations suited to the level and experience of your students. To select effectively, and to anticipate student questions, your must understand concepts at a higher level than you'll likely ever teach. Not all of our topics are central to the grade-school curriculum, and it is not our goal to teach you *how* to teach, though the ideas and approaches we'll explore are often suitable for a grade-school audience. The *mathematics* in this course shouldn't present much difficulty for math majors, requiring at most elementary calculus and a tiny bit of linear algebra; you should instead be considering how to *explain* the material, particularly to students with less mathematical knowledge than yourself.

We start with two motivational problems.[1]

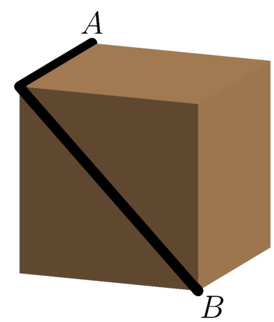

1. You wish to travel across the surface of a cube between two opposite vertices so that your path is as short as possible.

   Should you follow the path indicated?

   If yes, explain why.

   If not, how should you find the shortest path?

2. Two houses are to be connected to the electricity supply using a single connection.

   How should we determine where to place the connection so as to minimize the required length of wire?

   What information do you need in order to find the connection point?

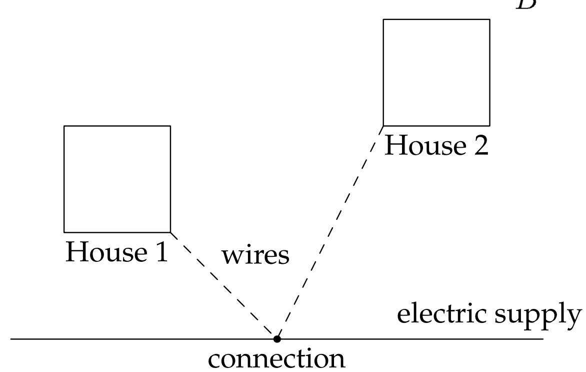

Your goal shouldn't only be to find the right answer! Consider how you might discuss these problems with grade-school students of different ability levels. Why might calculus *not* be a sensible approach? Are there any similarities between the two problems? Brainstorm some strategies...

[1]We are grateful to materials from UT Austin's UTeach program for suggesting several of the examples in this course including these motivational problems.

1 Sets & Functions

1.1 Basic Definitions

Consider how central functions are to mathematics, and how long you've been using them. How would you *define* "function" to someone with limited mathematical knowledge? Would you use words like *rule, assign, element, domain, vertical line test,* etc.? How helpful are these to your audience?

Examples 1.1. How would you explain the idea that the following do or do not represent functions?

1. $y = x^2$
2. Mon: fish, Tue: pork, Wed: fajitas, Thur: carbonara, Fri: pizza, Sat: fish, Sun: pizza
3. $(3,5)$, $(2,6)$, $(4,2)$, $(3,1)$.
4. $x^2 = y^2$

After considering the examples, perhaps you settle on a semi-formal definition:

A function f is rule which assigns to each input x exactly one output $f(x)$

Is this a useful definition? In what ways is it imprecise? Does the imprecision matter?

Of course the answers to these questions depend on your audience! What ideas do you want to convey to your students and can you do so without overburdening and intimidating them? To begin working towards a more complete picture, consider what we might allow to be *inputs* and *outputs.* This requires a small amount of set notation.

Definition 1.2. A *set* A is a collection of objects, or *elements.*[2] The notation $a \in A$ means that a is an element of A, sometimes read 'a lies in A.' Sets are often written upper case and elements lower.

A set B is a *subset* of a set A, written $B \subseteq A$, if every element of B is also an element of A: that is,

$$b \in B \implies b \in A$$

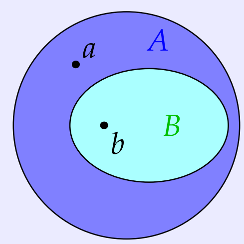

The picture illustrates sets A, B and elements a, b for which $B \subseteq A$, $a \in A$, $b \in B$ and $a \notin B$ (a does not lie in B).

Examples 1.3. 1. Suppose the elements of a set A are the numbers 1, 3, 5, 7 and 9. The simplest way to write this is using *roster notation*: we list the elements (in any order) between braces

$$A = \{1,3,5,7,9\}$$

Subsets are commonly expressed using *set-builder notation*. For example, here is a subset of A:

$$B = \{a \in A : 2 < a < 8\}$$

This is read, "The set of a in A such that a lies strictly between 2 and 8." In roster notation, $B = \{3,5,7\}$. Can you express B in other ways using set-builder notation?

[2] This is enough for our purposes, though a course in set theory will convince you that this definition has its own problems. Selection is always at work...

2. We summarize several common sets of numbers using informal combinations of roster and set-builder notation, all of which should be familiar.

   Natural numbers $\mathbb{N} = \{1, 2, 3, 4, \ldots\}$. For instance, $5 \in \mathbb{N}$ but $-3 \notin \mathbb{N}$.

   Integers $\mathbb{Z} = \{\ldots, -2, -1, 0, 1, 2, 3, \ldots\}$. For instance, $-4 \in \mathbb{Z}$ but $\frac{4}{5} \notin \mathbb{Z}$.

   Rational numbers (fractions) $\mathbb{Q} = \left\{\frac{p}{q} : p \in \mathbb{Z},\ q \in \mathbb{N}\right\}$. For instance $-\frac{6}{7} \in \mathbb{Q}$; in this case $p = -6$ is an integer, and $q = 7$ a natural number.

   Real numbers $\mathbb{R}$. For instance, $\sqrt{2} \in \mathbb{R}$. A formal definition is difficult, though we often informally visualize $\mathbb{R}$ as a *ruler*. *Intervals* are particularly important subsets, e.g.,

   $$[-4, \pi) = \{x \in \mathbb{R} : -4 \le x < \pi\}$$

   is a half-open interval.

   You should also be familiar with the *Cartesian plane*: $\mathbb{R}^2 = \{(x, y) : x, y \in \mathbb{R}\}$. The notation $(3, 4) \in \mathbb{R}^2$ here describes a *point* in the plane with *co-ordinates* $x = 3$, $y = 4$; don't confuse this with the *interval* $(3, 4) = \{x \in \mathbb{R} : 3 < x < 4\}$ which is a subset of $\mathbb{R}$!

   The subset relationships between these sets are in the order listed:

   $$\mathbb{N} \subseteq \mathbb{Z} \subseteq \mathbb{Q} \subseteq \mathbb{R}$$

   You should also have informally encountered the notion of *irrationality*: for instance, $\sqrt{2}$ and π are real numbers but not rational numbers.

The reason we need this language when discussing functions is that the inputs and outputs of a function are *elements* of sets. Here is a *very* formal definition of "function."

Definition 1.4. The *Cartesian product* of sets A, B is the set of *ordered pairs*

$$A \times B = \{(a, b) : a \in A, b \in B\}$$

A *function* from A to B is a non-empty subset $f \subseteq A \times B$ which satisfies the *vertical line test*

$$\text{For each } a \in A \text{, there is a } \textit{unique } b \in B \text{ such that } (a, b) \in f \qquad (*)$$

Instead of writing $f \subseteq A \times B$ and $(a, b) \in f$, we use the more familiar notation

$$f : A \to B \quad \text{and} \quad f(a) = b$$

To a function $f : A \to B$ are associated three useful sets:

- Domain: $\operatorname{dom} f = A$ is the set of *inputs.*
- Codomain: $\operatorname{codom} f = B$ is the set of *possible outputs.*
- Range: $\operatorname{range} f = \{b \in B : b = f(a) \text{ for some } a \in A\}$ is the set of *realized outputs.*

This probably isn't the definition you should give to 10^{th} graders, or even to freshman calculus students! But what should you do? How much of this is helpful in a a given context?

Example (1.1.2 cont.). We revisit our food-based example in this formal setting. To properly view this as a function $f : A \to B$, we have to carefully label the constituent sets.

$$A = \{\text{Mon, Tue, Wed, Thu, Fri}\}, \qquad B = \{\text{carbonara, fajitas, fish, pizza, pork}\},$$

$$f = \{(\text{Mon, fish}),\ (\text{Tue, pork}),\ (\text{Wed, fajitas}),\ (\text{Thu, carbonara}),\ (\text{Fri, pizza}),\ (\text{Sat, fish}),\ (\text{Sun, pizza})\}$$

The domain A should be clear, but we had to make a choice for the codomain B: in this case we chose it to equal to *range*. Can you suggest a different choice for B? Try the other examples yourself.

Representing Functions

Functions can be represented in various ways. We illustrate a few in an example.

Example 1.5. We consider the familiar *formula/rule* $f(x) = x^2$ in several contexts.

Table This presentation is most helpful when the domain is very small. The table shows the situation when $\operatorname{dom} f = \{-1, 0, 1, 2, 3\}$ and $\operatorname{range} f = \{0, 1, 4, 9\}$

x	-1	0	1	2	3
$f(x)$	1	0	1	4	9

Arrows A pictorial arrow diagram might also be helpful when the domain is small.

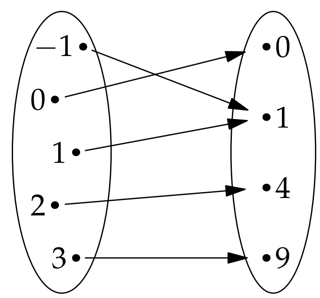

Graph This is the set of ordered pairs $\{(x, f(x)) : x \in \operatorname{dom} f\}$: in the context of the formal definition (1.4), *the graph is the function!*

For formulæ whose inputs and outputs are real numbers, two conventions are often observed:

- The domain is *implied* to be all real numbers for which the formula makes sense.
- The codomain is taken to be the set of real numbers.

If no other information is provided, we'd assume that the function defined by the formula $f(x) = x^2$ has both domain and codomain the entire set of real numbers: $f : \mathbb{R} \to \mathbb{R}$.

The range of the function is the set of possible outputs, in this case

$$\operatorname{range} f = \{x^2 \in \mathbb{R} : x \in \mathbb{R}\} = [0, \infty)$$

is the half-open interval of non-negative real numbers.

For 'calculus' functions like these, the vertical line test $(*)$ really involves vertical lines; every vertical line intersects the graph in precisely one point.

In the picture, the dots are the graph when the domain is the finite set $\{-1, 0, 1, 2, 3\}$ (as described in the table/arrow-diagram).

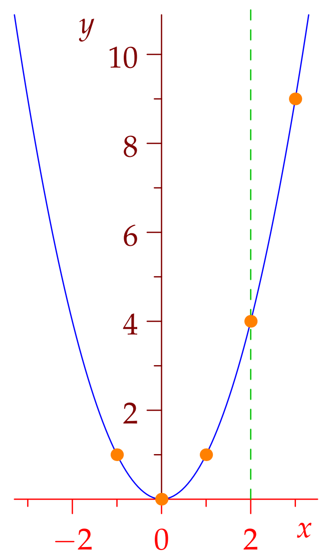

Can you think of other ways to represent a function? How might you decide which to use?

Exercises 1.1. 1. Let d represent the cost in millions of dollars to produce n cars, where n is measured in 1000s. As clearly as you can, explain what is meant by $d(25) = 431$.

2. A movie theater seats 200 people. For any particular show, the amount of money the theater takes in is a function of the number of people n in attendance. If a ticket costs $25, describe the domain and range of the function using set notation.

3. Temperature readings T were recorded every two hours from midnight to noon. Time t was measured in hours from midnight.

t	0	2	4	6	8	10	12
T (°F)	82	75	74	75	84	90	93

   (a) Plot the readings and use them to sketch a rough graph of T as a function of t.

   (b) Use your graph to estimate the temperature at 10:30 a.m.

4. State parts 1, 3 and 4 of Example 1.1 using the formal language of Definition 1.4. If you have a function, state the domain and range and explain how you know you have a function. If you don't have a function, explain why not.

   (*Since insufficient information is provided, there is no single correct answer!*)

5. (a) Let $A = \{1, 3, 5, 7, 9\}$. Explain in words what is meant by the set

   $$B = \{x \in A : x^2 > 10\}$$

   and state B in roster notation.

   (b) Find the set $C = \{x \in \mathbb{N} : (x-1)^2 < 16\}$ in roster notation.

   (c) Find the Cartesian product $B \times C$ in roster notation. Is it the same as $C \times B$?

6. Suppose that $f : \{-2, -1, 0, 1, 2\} \to \mathbb{R}$ is defined by the formula $f(x) = x^3 - 4x + 1$.

   Describe f using a table, an arrow diagram and a graph.

7. Find the implied domain and range for the functions defined by each rule:

   (a) $f(x) = \frac{x^2-4}{x-2}$ (b) $g(x) = \sqrt{x^2 - 16x}$ (c) $h(x) = \frac{1}{x}\sqrt{4x - x^2}$

   (*What is the largest set of real numbers for which the formula makes sense*?)

8. You ask your students to determine the range of the function f defined by the rule $f(x) = x^2$ with domain the interval $[-5, 2]$. You obtain various responses, including $[25, 4]$, $[4, 25]$, and $[-25, 4]$. What is going wrong? What is the correct answer, and how would you explain it to your students?

   More generally, if $\operatorname{dom} f = [a, b]$ (where $a \le b$), what is range f?

9. The unit circle is often represented by the implicit equation $x^2 + y^2 = 1$.

   (a) Draw the circle and explain why the full circle isn't the graph of a function.

   (b) Describe *two* functions $f : [-1, 1] \to \mathbb{R}$ and $g : [-1, 1] \to \mathbb{R}$ whose graphs together comprise the circle. What are the ranges of each function?

1.2 Linear Polynomials

Perhaps the simplest functions are the *linear polynomials*, whose graphs are straight lines,

$$y = f(x) = mx + c \quad \text{where } m, c \text{ are constants} \tag{$*$}$$

Linear polynomials make very simple models: increase the input by Δx and the output changes by $\Delta y = m\Delta x$ *regardless* of the starting value x. Given experimental data or a physical situation relating two quantities x and y, a *linear model* is an linear polynomial $(*)$ relating these variables. In practice, models are *approximations* to the real-world data. Later in the course we'll consider what should be meant by, and how to find, a 'good' linear model for approximately linear data.

Some of your earliest forays into algebra likely involved finding equations of straight lines.

Example 1.6. Find the equation of the straight line through the points $A = (1,3)$ and $B = (4,1)$.

Suppose the polynomial is $y = mx + c$. Since both A and B satisfy this equation, we start by substituting both points into the equation to find two relationships between m and c

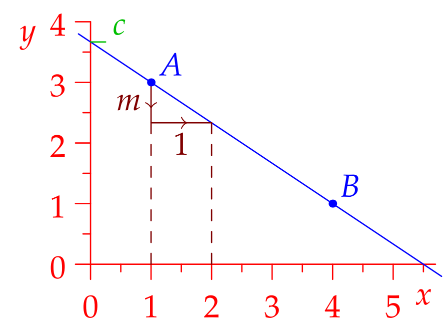

$$\begin{cases} 3 = m + c \\ 1 = 4m + c \end{cases}$$

This is a system of two linear equations in two unknowns (m, c). By now you should know several ways to solve such, but consider what might be easiest for a grade-school student...

Regardless of how you phrase it (solve one equation for c and substitute into the other, subtract one question from the other, etc.), we obtain

$$-2 = 3m \implies m = -\frac{2}{3} \implies c = 3 - m = \frac{11}{3}$$

whence the required polynomial is $y = \frac{1}{3}(11 - 2x)$.

As the picture suggests, the *gradient/slope* m represents how far one climbs/falls on travelling one unit to the right. The *y-intercept* c is the intersection of the graph with the vertical axis.

The above process works for any two points $A = (x_0, y_0)$ and $B = (x_1, y_1)$ provided $x_0 \neq x_1$: is it clear why this should be the case? The details are in Exercise 5. You might feel that such a problem is too abstract for your students, that such a 'proof' might be too intimidating. Indeed it might be counterproductive for some students, but consider several counterpoints:

- Once a student has developed comfort with concrete examples as above, Exercise 5 helps summarize and unify what they've learned. A general/abstract discussion helps build confidence by convincing a student that any such problem can be solved the same way.
- The most helpful elementary proofs are those which essentially replicate an example abstractly. Exercise 5 is not some abstract existence proof—it involves no trickery—it simply *reinforces* the core technique by applying it in the most general situation.
- Helping and encouraging students to think abstractly is one of the overarching learning outcomes of all mathematics. You might get push-back, but it's part of the job...

Example 1.7. Often the challenge of modeling lies in converting a word problem into algebra—don't underestimate how hard students find this! Here is a simple, though disguised, straight line model.

Beaker A contains a 300 ml solution of 2% acid. Beaker B contains 400 ml of acid of unknown concentration. The beakers are mixed together to produce an acid with concentration 6%. What was the concentration in beaker B?

Given your mathematical experience, it should seem natural to denote the unknown concentration (beaker B) by x. After mixing, we have a 700 ml solution containing $300 \times \frac{2}{100} + 400x$ ml of pure acid, whence its concentration is a linear polynomial function of x:

$$C(x) = \frac{6 + 400x}{700}$$

The problem is now easily solved: $C(x) = \frac{6}{100} \iff x = \frac{9}{100} = 9\%$.

Parametrized Lines Straight lines admit an alternative visualization. Imagine placing a *ruler* so that its zero point is at the origin $O = (0,0)$ and the "1" lies at a point $C = (c_1, c_2)$. If t (a real number) is the measure on the ruler, then the points on the line have co-ordinates

$$tC = (tc_1, tc_2) \qquad (*)$$

To describe the line through points A and B, place a ruler so that 0 corresponds to A and 1 to B. Now slide the ruler so that A moves to the origin O: this amounts to *subtracting* the co-ordinates of A from all points on the line. We obtain a parallel line through the origin, with B transformed to the point $C = B - A$. Putting this together with $(*)$ results in a parametrized description of the line:

$$(x, y) = A + tC = A + t(B - A) = (1 - t)A + tB$$

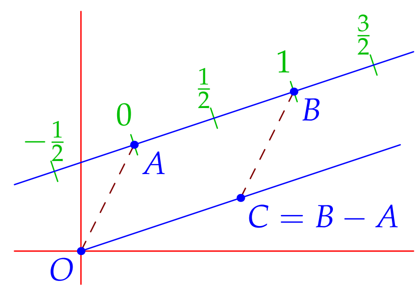

Contrast the parametrized description of a line with the linear polynomial approach: for instance, one challenge is that a line may be parametrized using infinitely many distinct rulers (choose *any* two points on the line!), whereas the linear polynomial description is unique. Does the parametrized approach have any advantages? Which description is easier to understand or to work with? Which fits better with your intuitive understanding of *line*? Which might cause a grade-school student the greater challenge?

In the Exercises we make sure that the two descriptions of a line correspond. The discussion is little more than the generalization of an example.

Example 1.8. The line through points $A = (3, 6)$ and $B = (-1, 4)$ may be parametrized by

$$(x, y) = (1 - t)(3, 6) + t(-1, 4) = \big(3 - 4t, 6 - 2t\big)$$

To convert this to a linear polynomial, first solve for t in terms of x,

$$x = 3 - 4t \implies t = \frac{1}{4}(3 - x)$$

before substituting into our expression for y:

$$y = 6 - 2t = 6 - \frac{2}{4}(3 - x) = \frac{1}{2}x - \frac{9}{2}$$

Exercises 1.2. 1. The cost of gasoline is \$4.20 per gallon on January 1st and \$4.90 on March 1st. State a *linear* function/model for the cost of gasoline as a function of time.

2. You have a choice of three different cell-phone plans.

   (a) No monthly charge and 10¢ per minute for all calls.

   (b) \$10 per month and 5¢ per minute for all calls.

   (c) \$30 per month, regardless of how many calls you make.

   How should you determine which plan to purchase?

3. Revisit Exercise 1.1.3. Find an approximate linear model $T(t) = mt + c$ for this data.

   (*There is no perfect answer*)

4. Revisit the beakers problem (Example 1.7). This time suppose we know that the concentration in beaker B is 9%. How much from beaker B should we pour into beaker A to obtain an acid with concentration 5%? Would you consider this a linear polynomial problem? Why/why not?

5. Suppose points $A = (x_0, y_0)$ and $B = (x_1, y_1)$ are given.

   (a) If $x_1 \neq x_0$, use the method of Example 1.6 to find the equation $y = mx + c$ of the line through these points.

   (b) Now use the parametrized approach where A corresponds to 0 and B to 1. If, in addition, $x_1 \neq x_0$, make things match up with your answer to part (a).
   What parametrization do you get if $A = (0, c)$ and $B = (1, m + c)$?

   (c) Part (a) provides an *algebraic* justification of the claim made on page 7, that the linear polynomial description of a line is unique ('*the* equation'). How might you help a student believe this claim if the algebra is unconvincing or too intimidating?
   (*Think about Example 1.6*)

6. A straight line is sometimes described as the set of points $(x, y) \in \mathbb{R}^2$ satisfying an equation of the form

   $$ax + by = c$$

   for some constants a, b, c where a, b are not both zero. How does this approach differ from our use of linear polynomials?

7. Throughout mathematics (particularly within *linear algebra*), a function $f : \mathbb{R} \to \mathbb{R}$ is said to be *linear* if it satisfies the condition

   $$\text{For all } \lambda, x \in \mathbb{R}, \quad f(\lambda x) = \lambda f(x)$$

   Is this the same thing as a linear polynomial? Explain.

1.3 Quadratic Polynomials

Quadratic polynomials are functions of the form $y = f(x) = ax^2 + bx + c$ where $a \neq 0$. The simplest is $y = x^2$, the standard parabola opening upwards. Here are some commonly encountered activities:

1. Find the *roots/zeros* of f, the solutions x to the equation $f(x) = 0$.
2. Sketch the *graph* of the function f.
3. Use quadratic functions to model a real-world problem.

You likely know two methods for finding zeros: factorizing and the quadratic formula, each of which has its problems. With experience it is easy to spot that

$$x^2 + 2x - 15 = (x-3)(x+5) = 0 \iff x = 3 \text{ or } x = -5$$

though the required creativity can make this difficult, particularly when coefficients are large. Students often prefer the quadratic formula since it always works, though at the cost of some intimidating algebra. We'll think about factorization shortly. First, we see how *completing the square* lies behind both the quadratic formula and the standard approach to graphing quadratic functions.

Example 1.9. Describe/graph the parabola $y = -3x^2 + 12x + 4$.

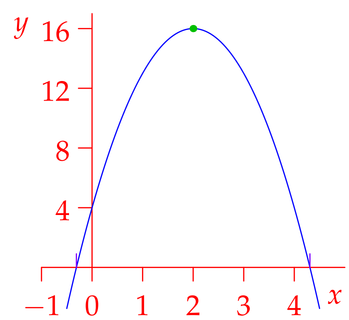

Pay attention to the x terms; $-3x^2 + 12x = -3(x^2 - 4x)$. Now

$$-3(x-2)^2 = -3(x^2 - 4x + 4) = -3x^2 + 12x - 12$$

gives most of what we want: note how we *divided the x-coefficient by two.* To finish, just tidy everything up,

$$y = (-3x^2 + 12x - 12) + 16 = -3(x-2)^2 + 16$$

The parabola therefore opens downwards ($-3 < 0$) with its apex (maximum) at $(x, y) = (2, 16)$.

This is easy, if intimidating, to repeat in general:

$$\begin{aligned} ax^2 + bx + c &= a\left(x^2 + \frac{b}{a}x\right) + c = a\left[\left(x + \frac{b}{2a}\right)^2 - \frac{b^2}{4a^2}\right] + c \\ &= a\left(x + \frac{b}{2a}\right)^2 - \frac{b^2 - 4ac}{4a} \qquad (*) \end{aligned}$$

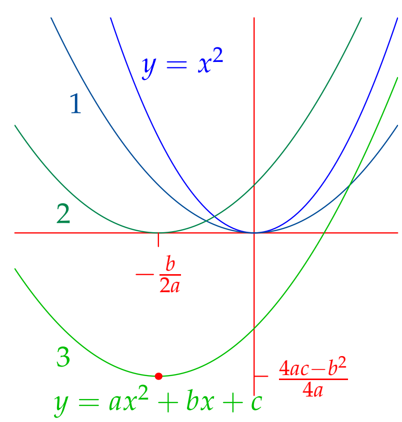

The graph is that of the standard parabola which has been:

1. Vertically scaled by a;
2. Shifted horizontally by $-\frac{b}{2a}$;
3. Shifted vertically by $\frac{4ac-b^2}{4a}$

By solving $(*)$ for x, we see that completing the square yields the *quadratic formula.*

Theorem 1.10. *If $a \neq 0$, then* $ax^2 + bx + c = 0 \iff x = \dfrac{-b \pm \sqrt{b^2 - 4ac}}{2a}$

Example (1.9 cont). Our analysis suggests two methods for finding the roots.

1. Quadratic formula: with $a = -3$, $b = 12$, $c = 4$, we have

$$x = \frac{-12 \pm \sqrt{12^2 - 4(-3)\cdot 4}}{2(-3)} = \frac{-12 \pm 4\sqrt{3^2+3}}{-6} = 2 \pm \frac{\sqrt{12}}{3} = 2 \pm \frac{2\sqrt{3}}{3}$$

While it is always tempting to jump for a formula, it often leads to difficult surd expressions. We simplified by noticing the common factor of 4^2 inside the square root. Without this, we'd be faced with $\sqrt{144+48} = \sqrt{192}$.

2. Use the fact that we've already completed the square:

$$-3(x-2)^2 + 16 = 0 \iff (x-2)^2 = \frac{16}{3} \iff x = 2 \pm \frac{4}{\sqrt{3}}$$

In many cases it is simpler to complete the square than to use the quadratic formula—remember that they are equivalent!

Polynomials are often employed in modelling due to their simplicity and ease of evaluation. As you saw in calculus, the motion of a falling body, or of any projectile can be modelled using quadratic polynomials, an observation going back to at least to Galileo in the early 1600s: the distance travelled by a falling body is proportional to the *square* of the time taken $y(t) - y(0) \propto t^2$.

Example 1.11. A body is dropped from a height of 125 meters, taking exactly 5 seconds to reach the ground. Its height at time t seconds is given by $y(t) = 125 - 5t^2$ m.

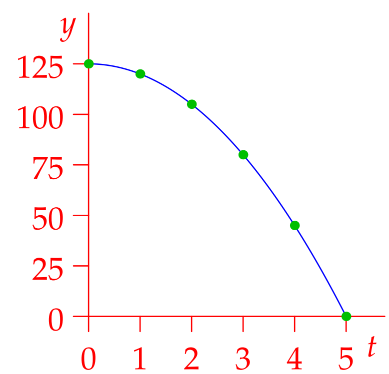

This certainly fits Galileo's observation: $y(t) - y(0) = -5t^2$ is indeed proportional to t^2.

Over each interval of 1 s, we may ask how far the body falls; we summarize in a table.

t	0	1	2	3	4	5
$y(t)$	125	120	105	80	45	0
$y(t) - y(0)$	0	−5	−20	−45	−80	−125
Δy	−5	−15	−25	−35	−45	

Since each interval has duration 1 s, each Δy is the *average speed* of the falling body over that interval.

You'll have seen problems like this in calculus; likely you want to *differentiate* to find the *velocity* $y'(t) = -10t$ m/s and *acceleration* $y''(t) = -10\,\text{m/s}^2$. However, historically and in introductory calculus, it is problems like these that *motivate the definition* of the derivative.[3]

Armed with calculus, Galileo's observation is that the height $y(t)$ solves a differential equation

$$\frac{d^2y}{dt^2} = -g \implies y'(t) = -gt + v_0 \implies y(t) = -\frac{1}{2}gt^2 + v_0 t + h_0$$

where g (approximately $32\,\text{ft/s}^2$ or $10\,\text{m/s}^2$) is the constant acceleration due to gravity, and the constants of integration h_0, v_0 are the initial height and vertical velocity. Unless you are explicitly teaching calculus or Newtonian physics, this is probably a bad place to start!

[3] The last line of the table really does suggest that speed is a linear function!

Example 1.12. Your frisbee is stuck 15 m up a tree. Standing 10 m from the base of the trunk, you throw a ball with the intent of knocking the frisbee out of the tree.

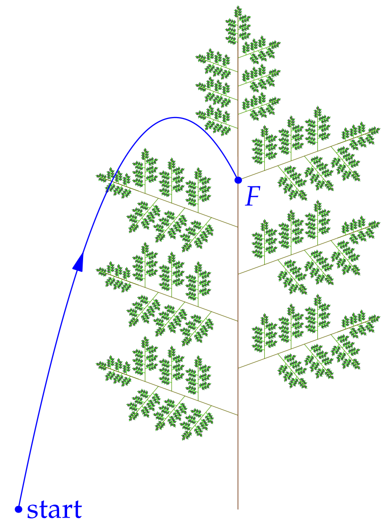

The standard approach to modeling such problems involves considering the horizontal and vertical motions separately.

Horizontal $x(t) = pt + q$ is a *linear function* of time.

Vertical $y(t) = -5t^2 + rt + s$ is a *quadratic function* of time.

Substituting for t yields a quadratic function for the trajectory

$$y(x) = ax^2 + bx + c$$

We'll leave the details of the solution to Exercise 6. For the present, consider why there are *multiple answers*; can you explain why *without* explicitly solving the problem?

Exercises 1.3. 1. Complete the square for each quadratic function. Use your answer to find the range and to graph the function.

(a) $f(x) = x^2 - 6x + 5$ (b) $f(x) = -x^2 + x + 1$

(c) $f(x) = -3x^2 + 8x + 5$

2. For the quadratic function $y = 2x^2 - 5x + 7$, produce a table for $x \in \{0, 1, 2, 3, 4, 5, 6\}$ similarly to that in Example 1.11. What do you observe about Δy?

3. Find the implied domain of the function $f(x) = \frac{1}{\sqrt{4-7x+x^2}}$

4. (a) Find the equations of all quadratic polynomial functions which pass through the points $(1, 3)$ and $(2, 4)$.

   (b) More generally, if $P = (a, b)$ and $Q = (c, d)$ are given, where $c \neq a$, find all quadratic functions whose graphs contain P and Q.

5. Describe as best you can how the graph of the function $f(x) = 3x^2 + bx + 2$ depends on b.

6. Consider the frisbee/tree problem (Example 1.12). Assume you're standing at the origin and that the frisbee is at the point $(10, 15)$.

   (a) Find/describe all suitable trajectories that result in the ball hitting the frisbee.

   (b) (Hard) Find a formula relating the initial speed v and initial slope m of the parabola (the initial speed/direction in which you throw the ball).

      i. If you throw the ball in such a way that the initial *vertical* speed of the ball is twice its *horizontal* speed, find how fast you have to throw the ball in order to hit the frisbee.

      ii. What is the *minimum* speed at which you could throw the ball if you want to dislodge the frisbee?

   (*Hint: You'll need some calculus! In the language of the original problem, the initial slope is* $m = \frac{r}{p}$ *and the initial speed* $v = \sqrt{p^2 + r^2}$; *why?*)

1.4 Polynomials, Factorization & the Rational Roots Theorem

Recall our simple example of factorization in the previous section

$$x^2 + 2x - 15 = (x-3)(x+5) = 0 \iff x = 3 \text{ or } x = -5$$

That this approach provides *all* roots relies on several familiar algebraic facts:

1. Factor Theorem: $f(c) = 0 \iff x - c$ is a *factor* of $f(x)$.
2. No zero-divisors: $g(x)h(x) = 0 \iff g(x) = 0$ or $h(x) = 0$.
3. A quadratic has *at most two* distinct roots.

We'll examine this more closely at the end of this section. For students first learning factorization, it isn't the *why* that's the challenge, it's the *how*. Multiplying out $(x-3)(x+5)$ is mechanical, but factorizing requires some creativity; we can't really factor without somehow knowing that 3 and -5 are roots! Beyond making a lucky guess, how might we go about this?

Example 1.13. Let's re-examine $f(x) = x^2 + 2x - 15 = 0$ in a couple of stages.

Integer solutions The simplest type of root would be an *integer* n. If $f(n) = 0$, observe that

$$n^2 + 2n - 15 = 0 \implies n(n+2) = 15 \implies 15 \text{ is divisible by } n$$

There are only *eight possible candidates* for n, and it doesn't take long to test them all:

n	1	-1	3	-3	5	-5	15	-15
$n+2$	3	1	5	-1	7	-3	17	-13

Rather than computing $f(n)$ explicitly, we listed all divisors of n in the first, the corresponding $n+2$ in the second, and mentally checked when $n(n+2) = 15$. There are precisely two integer solutions, namely $n = 3$ and $n = -5$.

Rational Solutions If you already believe that a quadratic polynomial has *at most* two solutions, then you're done. The next simplest possibility, however, is that a solution be a *rational number* $x = \frac{p}{q}$: we may assume this is in *simplest terms*.[4] Substituting into the polynomial, we see that

$$\frac{p^2}{q^2} + 2\frac{p}{q} - 15 = 0 \iff p^2 + 2pq - 15q^2 = 0$$

Remembering that p, q are *integers*, we rearrange this equation in two ways:

$p(p+2q) = 15q^2$ Since the left side is a multiple of p, so also is the *right*. Since p, q have no common factors, it follows that p divides into 15 (15 is a multiple of p).

$p^2 = q(15q - 2p)$ Since the right side is a multiple of q, so also is the *left*. Since p, q have no common factors, we conclude that $q = 1$.

The upshot is that the only rational solutions to $f(x) = 0$ are the two *integers* we've already found.

[4] I.e., $p \in \mathbb{Z}$ and $q \in \mathbb{N}$ have no common factors: $\gcd(p, q) = 1$.

Definition 1.14. A *degree n polynomial* is any function of the form

$$f(x) = a_n x^n + a_{n-1}x^{n-1} + \cdots + a_1 x + a_0$$

where the *coefficients* a_k are constants with $a_n \neq 0$.

A quadratic polynomial has degree 2 and a linear polynomial $mx + c$ degree one[5] (if $m \neq 0$). Our analysis in Example 1.13 generalizes to a famous result.

Theorem 1.15 (Rational Roots). *Suppose* $f(x) = a_n x^n + \cdots + a_0$ *has* integer *coefficients where* a_n *and* a_0 *are non-zero. If* $x = \frac{p}{q}$ *is a rational root in simplest terms, then* q *divides into* a_n *and* p *into* a_0*. In particular, if* $a_n = 1$*, then the only possible rational roots are integers.*

Proof. Substitute $\frac{p}{q}$ into $f(x)$ and multiply by q^n to obtain an equation where everything is an *integer*

$$\underbrace{a_n p^n + \overbrace{a_{n-1}p^{n-1}q + \cdots + a_1 pq^{n-1}}^{\text{divisible by } q}}_{\text{divisible by } p} + a_0 q^n = 0$$

By considering the braced terms we see that $a_n p^n$ is divisible by q and $a_0 q^n$ by p. Since p, q have no common factors, we obtain the result. ∎

Examples 1.16. 1. If $x = \frac{p}{q}$ is a rational root of $f(x) = 2x^2 - x - 3$ in lowest terms, then $q = 1$ or 2 and $p = \pm 1$ or ± 3. The eight possibilities for x are easily checked:

x	1	-1	3	-3	$\frac{1}{2}$	$-\frac{1}{2}$	$\frac{3}{2}$	$-\frac{3}{2}$
$2x-1$	1	-3	5	-7	0	-2	2	-4

You may prefer to compute $f(x)$ directly: as in the previous example, since we already know x it is quicker to check whether $x(2x - 1) = 3$ rather than $f(x) = 0$ (consider whether this trick would be helpful or confusing in a grade-school context). The two roots are indicated; it is easily verified that the polynomial can be factorized $f(x) = (2x - 3)(x + 1)$.

2. If the cubic polynomial $f(x) = x^3 - 2x^2 + 5$ had any rational roots, the only possibilities would be ± 1 or ± 5. It is quickly verified that none of these work,

$$f(1) = 4, \quad f(-1) = 2, \quad f(5) = 80, \quad f(-5) = -170$$

whence $f(x) = 0$ has no rational roots.

Unless there are very few candidates for rational roots, checking all possibilities by hand is time-consuming. The rational roots theorem is therefore typically used in conjunction with factorization by providing options for how to start factorizing. This still isn't easy, as the next example shows.

[5] A non-zero constant polynomial has degree zero. By convention, the *zero polynomial* $y \equiv 0$ has degree $-\infty$ so that the theorem $\deg fg = \deg f + \deg g$ makes sense for all polynomials.

Example 1.17. Consider the cubic function $f(x) = x^3 - x^2 - 7x + 10$. The rational roots theorem offers eight candidates for rational roots: $x = \pm 1, \pm 2, \pm 5, \pm 10$. It is not difficult to check the first few of these in your head, for instance,

$$f(2) = 8 - 4 - 14 + 10 = 0$$

By the factor theorem, $x - 2$ is a factor of $f(x)$. The factorization can be performed in various ways. Here are three options, though all are versions of the same process.

Long/synthetic division You should have practiced this in high-school.

$$\begin{array}{r l} & x^2 + x - 5 \\ x-2 \,) & x^3 - x^2 - 7x + 10 \\ & -x^3 + 2x^2 \\ & \quad x^2 - 7x \\ & \quad -x^2 + 2x \\ & \qquad -5x + 10 \\ & \qquad 5x - 10 \\ & \qquad\qquad 0 \end{array} \implies x^3 - x^2 - 7x + 10 = (x-2)(x^2 + x - 5)$$

Multiply out and solve Write $f(x) = (x-2)q(x)$ where $q(x) = ax^2 + bx + c$ is some quadratic polynomial. Now multiply out:

$$x^3 - x^2 - 7x + 10 = (x-2)(ax^2 + bx + c) = ax^3 + (b - 2a)x^2 + (c - 2b)x - 2c$$

Equating coefficients, we obtain the same factorization as before:

$$a = 1, \qquad b - 2a = -1 \implies b = 1, \qquad -2c = 10 \implies c = -5$$

Term-by-term factorization We construct the required quadratic factor term-by-term. Since each calculation can be done in your head, with practice you'll find that you can factorize in one line without showing any work. *Teaching* such an approach is likely a terrible idea unless your students are already very comfortable with factorization!

(a) To create x^3, the first term of the quadratic factor must be x^2

$$x^3 - x^2 - 7x + 10 = (x-2)(x^2 + \cdots) = x^3 - 2x^2 + \cdots$$

(b) We have $-2x^2$ but want $-x^2$. To correct this, add x to the quadratic ($x^2 - 2x^2 = -x^2$):

$$(x-2)(x^2 + x + \cdots) = x^3 - x^2 - 2x + \cdots$$

(c) We have $-2x$ but want $-7x$. To fix, subtract 5 from the quadratic ($-5x - 2x = -7x$):

$$(x-2)(x^2 + x - 5) = x^3 - x^2 - 7x + 10$$

(d) Since the last term 10 is correct, the factorization worked!

You might have seen other approaches involving arranging the coefficients in a table. Regardless, the calculations required to complete these methods are exactly those seen above; all these methods are versions of the same thing.

Why Does Factorization Work?

The theory of factorization relies on some algebra. Here is a *brief* treatment.

Theorem 1.18 (Factor Theorem). *Suppose $f(x)$ is a degree n polynomial. Then:*

1. *$f(c) = 0$ if and only if $f(x) = (x-c)q(x)$ for some (degree $n-1$) polynomial $q(x)$.*
2. *The polynomial has* at most *n distinct roots.*

Proof. 1. ($\Leftarrow$) This is essentially trivial: $f(x) = (x-c)q(x) \implies f(c) = (c-c)q(c) = 0$.

($\Rightarrow$) This relies on the *division algorithm for polynomials*: if f, g are polynomials, then there are unique polynomials q, r with[6]

$$f(x) = g(x)q(x) + r(x) \quad \text{and} \quad \deg r < \deg g$$

If $g(x) = x - c$ is linear, $r(x)$ must be constant. Evaluate both sides at $x = c$ to obtain

$$f(x) = (x-c)q(x) + f(c) \qquad \text{(thus } f(c) = 0 \Longrightarrow f(x) = (x-c)q(x)\text{)}$$

2. Suppose $c_1, \ldots, c_n$ are distinct real roots. By part 1, $f(x) = (x-c_1)q_1(x)$. Since

$$0 = f(c_2) = (c_2 - c_1)q_1(c_2) \implies q_1(c_2) = 0$$

we may factor $x - c_2$ from $q_1(x)$ to obtain

$$f(x) = (x-c_1)(x-c_2)q_2(x), \qquad \deg q_2 = n-2$$

Repeat this process to factor out all n linear polynomials $x - c_k$:

$$f(x) = (x-c_1)\cdots(x-c_n)q_n, \qquad \deg q_n = n - n = 0$$

whence $q_n \neq 0$ is *constant*. Plainly $f(c) = (c-c_1)\cdots(c-c_n)q_n = 0 \Longrightarrow c = c_j$ for some j, so there are no other roots. ■

Example (1.17 cont). We know that $f(x) = x^3 - x^2 - 7x + 10 = (x-2)(x^2+x-5)$. But then

$$f(x) = 0 \iff x - 2 = 0 \text{ or } x^2 + x - 5 = 0$$

The former gives the root $x = 2$, and the latter can be attacked via the quadratic formula or completing the square; the polynomial therefore has exactly three real roots

$$x = 2, \ \frac{-1 \pm \sqrt{21}}{2}$$

[6] For a given example, q and r may be found by synthetic division. This is similar (and may be demonstrated similarly) to the more familiar division algorithm for integers: if m, n are integers, then there are unique integers q, r for which

$$m = qn + r \text{ and } 0 \le r < |n|$$

In elementary school, this is typically written $m \div n = q \,\mathrm{r}\, r$ (q remainder r); e.g., $23 \div 4 = 5 \,\mathrm{r}\, 3$ corresponds to $23 = 5 \times 4 + 3$.

Example 1.19. We finish with a quick example of how long division (or any other factorization method as in Example 1.17) computes the ingredients in the division algorithm.

If $f(x) = x^3 + 7x^2 - 2$ and $g(x) = x^2 - 2$, then

$$
\begin{array}{r|l}
 & \qquad\qquad\quad x \quad +7 \\ \hline
x^2-2 & x^3+7x^2 \qquad\quad -2 \\
 & -x^3 \qquad\quad +2x \\ \hline
 & \qquad 7x^2+2x \quad -2 \\
 & \qquad -7x^2 \qquad\; +14 \\ \hline
 & \qquad\qquad\quad 2x+12
\end{array}
$$

$\implies x^3 + 7x^2 - 2 = (x^2 - 2)(x + 7) + (2x + 12)$

Otherwise said, $f(x) = g(x)q(x) + r(x)$, where

$q(x) = x + 7,\ r(x) = 2x + 12$ and $\deg r = 1 < 2 = \deg g$.

Exercises 1.4.

1. Apply the rational roots theorem to the polynomial $x^3 + 2x^2 - x - 2$ and use it to factorize the polynomial.
2. Repeat the previous question for the polynomial $6x^2 + x - 2$.
3. Use the rational roots theorem to prove that the polynomial $2x^5 - 3x + 7$ has no rational roots.
4. Factorize the polynomials and thereby find their (real) roots. Explain your steps carefully.
   (a) $f(x) = x^3 + 2x^2 - 3x$
   (b) $f(x) = x^4 - 13x^2 + 36$
   (c) $f(x) = x^3 - 7x - 6$
5. Factorize the polynomial $f(x) = x^6 - 2x^5 - x^4 - 4x^3 - 4x^2 - 4x - 6$ and thus demonstrate that it has exactly two real roots.
6. Students often follow a heuristic when trying to factorize a polynomial $f(x) = 0$: try some small integer values for x until you find a root, then apply long division. For what types of polynomial $f(x)$ will this approach work? Explain.
7. The polynomial $f(x) = 2x^4 - 3x^3 + 2x^2 + 3x - 9$ has only one rational root. Find it and factorize the polynomial as $f(x) = g(x)q(x)$ where $\deg g = 1$.
8. Find unique polynomials $q(x)$ and $r(x)$ for which $f(x) = g(x)q(x) + r(x)$ and $\deg r < \deg g$.
   (a) $f(x) = x^3 + 1$ and $g(x) = x + 2$.
   (b) $f(x) = x^4 + x^3 - 2$ and $g(x) = x^2 + 1$.
9. Let $f(x) = ax^3 + bx^2 + cx + d$ be a cubic polynomial. 'Complete the cube' by finding a constant k such that
   $$f(x) = a(x-k)^3 + p(x-k) + q$$
   has no $(x-k)^2$ term (here p, q are constants).
   (*Hint: evaluate* $f(x+k)$)
10. Suppose $\deg f = k$ and $\deg g = l$.
    (a) Show that $\deg(fg) = kl$.
    (b) Is it always the case that $\deg(f + g) = \max(k, l)$? Why/why not?

1.5 Inverse Functions & the Horizontal Line Test

The informal idea of an inverse function is that f^{-1} takes the *output* of f and returns its *input* (and vice versa).

Example 1.20. Define a simple function using a table or an arrow diagram

x	1	2	3	4
$f(x)$	4	2	5	7

y	4	2	5	7
$f^{-1}(y)$	1	2	3	4

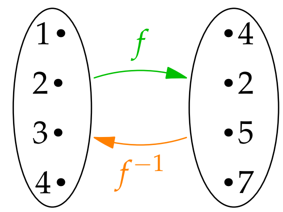

The inverse f^{-1} is the function obtained by *reversing the arrows* or flipping the table upside-down.

Definition 1.21. A function $f : A \to B$ is *invertible* if it has an *inverse*: a function $f^{-1} : B \to A$ for which

$$f^{-1}(f(x)) = x \quad \text{and} \quad f(f^{-1}(y)) = y \tag{$*$}$$

for all possible inputs $x \in A$ and $y \in B$.

Certainly Example 1.20 satisfies the input–output properties $(*)$. Our concerns are identifying when a function is invertible, how to make it so if not, and how to compute an inverse.

Examples 1.22. 1. The function $f(x) = 2x$ has inverse $f^{-1}(y) = \frac{y}{2}$.

The input–output conditions $(*)$ are certainly satisfied.

The graph admits an interpretation of f^{-1} similar to the arrow diagram.

- The function f takes an input x, moves it vertically to the graph, then projects to the y-axis. This interpretation is precisely the vertical line test (Definition 1.4)!
- The inverse function *reverses the arrows*: transport an input y horizontally to the graph, then project to the x-axis.

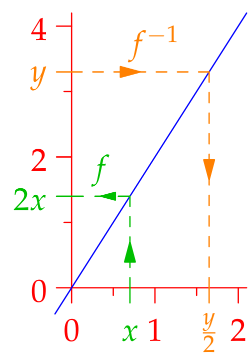

2. Consider $f(x) = x^2 - 1$. This time, when attempting to move a real number y horizontally to the graph, we usually encounter one of two problems:

   (a) If $y > -1$, there are two choices of x (two intersections).

   (b) If $y < -1$, there is no intersection with the graph.

   The naïve approach of *reversing the arrows* is insufficient to define an inverse. However, a simple remedy arises by staring at the graph:

   - Problem (a) goes away if we delete the left half of the graph. Equivalently, we *restrict the domain* of f to $[0, \infty)$.
   - Problem (b) disappears if we insist that $y \geq -1$. Equivalently, we *restrict the codomain* of f to its *range* $[-1, \infty)$.

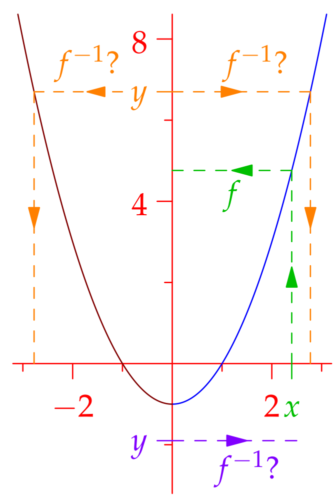

After making these restrictions so that $f : [0, \infty) \to [-1, \infty)$, it is easily checked that

$$f^{-1}(y) = \sqrt{y+1}, \quad f^{-1} : [-1, \infty) \to [0, \infty)$$

satisfies the input–output conditions $(*)$ and is therefore the inverse of f:

$$x \in [0, \infty) \implies f^{-1}(f(x)) = \sqrt{(x^2 - 1) + 1} = x$$
$$y \in [-1, \infty) \implies f(f^{-1}(y)) = \left(\sqrt{y+1}\right)^2 - 1 = y$$

What makes a function invertible? The fixes in the last example can be rephrased succinctly:

Horizontal line test: every horizontal line must intersect the graph *exactly once*

This unpacks to two conditions, each of which addresses one of the problems seen in the example.

Definition 1.23. Let $f : A \to B$ be a function. We say that f is:

(a) *1–1/one-to-one* if distinct inputs $x_1 \neq x_2 \in A$ have distinct outputs $f(x_1) \neq f(x_2)$. Equivalently,

Given $x_1, x_2 \in A$, we have $f(x_1) = f(x_2) \implies x_1 = x_2$

If A, B are sets of real numbers, each horizontal line intersects the graph at most once.

(b) *Onto* if range $f = B$. Equivalently,

Given $y \in B$, there is some $x \in A$ for which $y = f(x)$

If $A, B \subseteq \mathbb{R}$, the horizontal line through $y \in B$ intersects the graph at least once.

Putting these ideas together, a function is both 1–1 and onto precisely when every $y \in B$ corresponds to a *unique* $x \in A$ for which $y = f(x)$. In summary:

Theorem 1.24. *$f : A \to B$ is invertible if and only if it is both 1–1 and onto. Its inverse is the function $f^{-1} : B \to A$ such that $f^{-1}(y) = x$ whenever $y = f(x)$.*

Example (1.22.2, mk. II). Consider the two properties in the context of the example $f(x) = x^2 - 1$:

(a) $f(x_1) = f(x_2) \implies x_1^2 - 1 = x_2^2 - 1 \implies x_1^2 = x_2^2 \implies x_1 = \pm x_2$.
To force f to be 1–1, it is enough to *restrict the domain* so that all x have the same sign: the obvious choice is dom $f = [0, \infty)$.

(b) range $f = \{x^2 - 1 : x \in [0, \infty)\} = [-1, \infty)$. We force f to be onto by *restricting its codomain* to $[-1, \infty)$.

The inverse function is obtained by solving $y = x^2 - 1$ for x:

$$x^2 = y + 1 \implies x = f^{-1}(y) = \sqrt{y+1}$$

The *non-negative square root* is used since $x \in \operatorname{dom} f = [0, \infty)$.

An algorithm for inverting functions Our discussion provides an algorithmic process for making a function $f : A \to B$ invertible and finding an inverse.

(a) Check that f is 1–1. If not, *restrict the domain* until it is.

(b) Check that f is onto. If not, *redefine* $B = \operatorname{range} f$.

(c) Solve $y = f(x)$ for $x = f^{-1}(y)$.

Since x is typically preferred as an input, it is common to *switch* x, y at the end of step 3 and write $y = f^{-1}(x)$. If $A, B \subseteq \mathbb{R}$, switching $x \leftrightarrow y$ is equivalent to *reflecting the graph* in the line $y = x$.
Note also that step (a) likely involves a *choice*; depending on how you restrict the domain, you can find multiple inverse functions! To see this in action, we return once more to our example.

Example (1.22.2, mk. III). Recall that if $f(x) = x^2 - 1$, then

$$f(x_1) = f(x_2) \implies x_1 = \pm x_2$$

Instead of restricting the domain to $[0, \infty)$, we can instead force f to be 1–1 by taking the other half of the graph; by *choosing* $\operatorname{dom} f = (-\infty, 0]$. The range/codomain remains $[-1, \infty)$, but the inverse function is now different:

$$x^2 = y + 1 \implies x = -\sqrt{y+1} \in (-\infty, 0] = \operatorname{dom} f \implies f^{-1}(x) = -\sqrt{x+1}$$

This time the new domain for f forced us to use the *negative square root.*

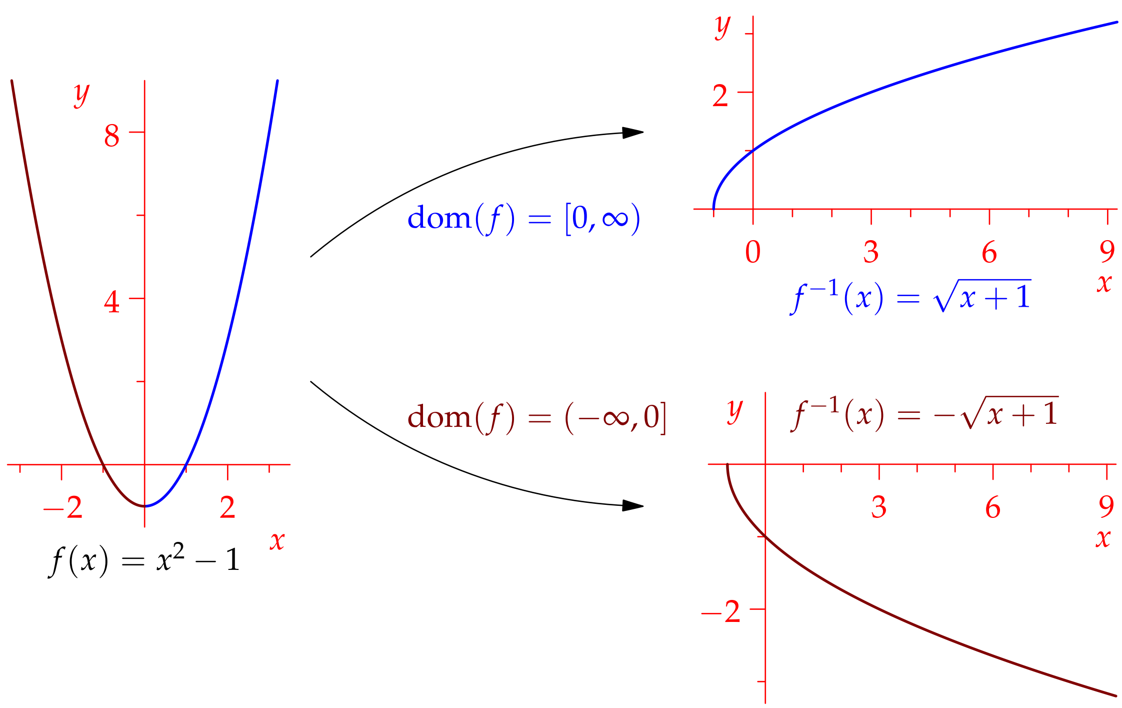

We could choose other domains on which f is 1–1, but these are the most natural choices.

The moral is that you cannot invert a function unless you are precise about its domain and range!

We finish with an algebraically tougher example, where you may feel that more detail is justified.

Example 1.25. Let $y = f(x) = \frac{1}{(x-2)^2}$. Its implied *domain* consists of all real numbers except 2.

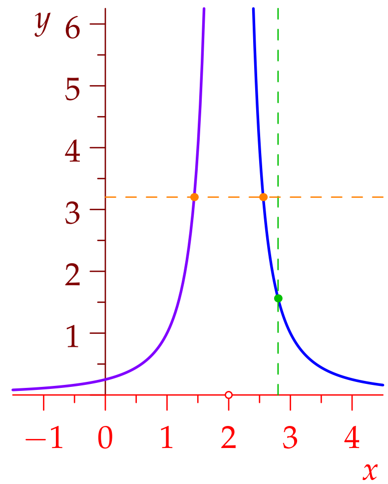

The *vertical line test* is clearly visible on the graph: every vertical line $x = a$, except $x = 2$, intersects the graph exactly once.

The *range* is the interval $\mathbb{R}^+ = (0, \infty)$ as can be seen by solving

$$f(x) = y \iff \frac{1}{x-2} = \pm\sqrt{y} \iff x = 2 \pm \frac{1}{\sqrt{y}}$$

Any positive output y may be obtained via $y = f\big(2 + \frac{1}{\sqrt{y}}\big)$.
The $\pm$-term shows that f fails the *horizontal line test*: it isn't 1–1.

There are two natural choices for an inverse:

(a) Choose $\operatorname{dom} f = (2, \infty)$, then $\pm\sqrt{y} = \frac{1}{x-2}$ is *positive*. We take the *positive* square root and obtain the inverse function

$$g : (0, \infty) \to (2, \infty), \quad g(x) = 2 + \frac{1}{\sqrt{y}}$$

(b) Choose $\operatorname{dom} f = (-\infty, 2)$, then $\pm\sqrt{y} = \frac{1}{x-2}$ is *negative* and we obtain a second inverse function

$$h : (0, \infty) \to (-\infty, 2), \quad h(x) = 2 - \frac{1}{\sqrt{y}}$$

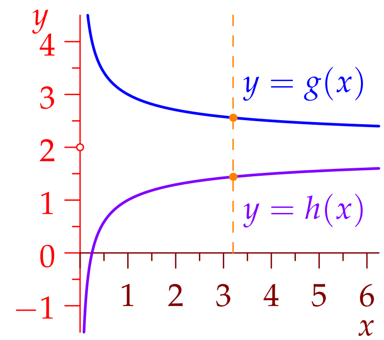

Exercises 1.5. 1. If $\operatorname{dom} f = \mathbb{R}$, check that $f(x) = x^3 + 8$ passes the horizontal line test. Find f^{-1}.

2. Consider $f(x) = x^2 + 2x - 3$. Similarly to Example 1.22, find *two inverses* of f.

3. Sketch the graph of the function

$$f(x) = \begin{cases} x & \text{if } 0 \le x < 1 \\ x - 1 & \text{if } 1 \le x < 2 \\ x - 2 & \text{if } 2 \le x < 3 \end{cases}$$

Find *three* domains on which f is 1–1 and thus compute three distinct inverses.

4. Show that the following function $f : \mathbb{R} \to (\frac{3}{2}, \infty)$ is 1–1 and onto, sketch its graph and find f^{-1}.

$$f(x) = \begin{cases} 3 - \frac{1}{2}x & \text{if } x \le 2 \\ 2 - \frac{1}{x} & \text{if } x > 2 \end{cases}$$

5. (Hard) Find the implied domain and range of $f(x) = \frac{x+1}{1+\frac{1}{x+1}}$. Now find an interval on which f is 1–1 and compute its inverse.

6. An astute student observes that Definition 1.21 only describes the properties satisfied by *an* inverse and asks why we keep referring to *the* inverse. How would you respond?

2 Trigonometric Functions and Polar Co-ordinates

In this chapter we review trigonometry and periodic functions and discuss their relation to polar co-ordinates. Some of this will be non-standard.

2.1 Definitions & Measuring Angles

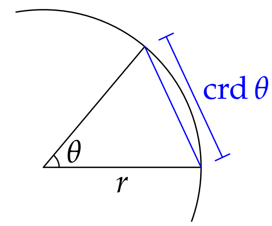

Trigonometric functions date back at least 2000 years. Ancient mathematicians were interested in the relationship between the *chord* of a circle and the central angle, often for the purpose of astronomical measurement. It wasn't until 1595 that the term *trigonometry* (literally *triangle measure*) was coined, and the functions were considered as coming from triangles.

Here are several related definitions of sine, cosine and tangent based either on triangles or circles.

Definition 2.1. 1. (a) Given a right triangle with *hypotenuse* (longest side) 1 and angle θ, define $\sin\theta$ and $\cos\theta$ to be the side lengths *opposite* and *adjacent* to θ.

Define $\tan\theta = \frac{\sin\theta}{\cos\theta}$ to be the slope of the hypotenuse.

(b) Given a right triangle with angle θ, *hypotenuse* r, *adjacent* x and *opposite* y, define

$$\sin\theta = \frac{y}{r} \qquad \cos\theta = \frac{x}{r} \qquad \tan\theta = \frac{y}{x}$$

2. (a) $(\cos\theta, \sin\theta)$ are the co-ordinates of a point on the unit circle, where θ is its *polar angle* measured counter-clockwise from the positive x-axis. Provided $\cos\theta \neq 0$, also define $\tan\theta = \frac{\sin\theta}{\cos\theta}$.

(b) Repeat the definition for a circle of radius r with co-ordinates $(r\cos\theta, r\sin\theta)$.

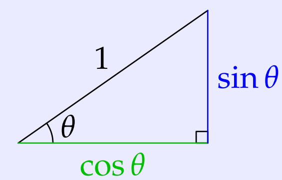

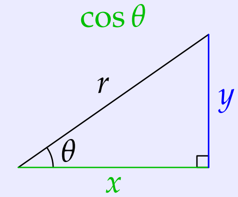

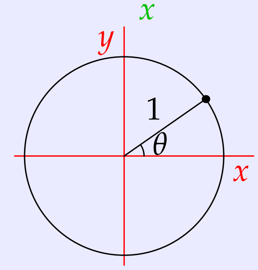

Discuss some of the advantages and weaknesses of these definitions:

- What prerequisites are you assuming in each case?
- Is it easier to think about *lengths* rather than ratios?
- Where do you need basic facts from Euclidean geometry such as *congruent/similar* triangles?
- Convince yourself that that the triangle definitions follow from the circle definitions. What is missing if you try to use the triangle definition to justify the circle version?
- If you were introducing trigonometry for the first time, what would you use?

If you've done sufficient calculus you might know of other definitions, for instance using power (Maclaurin) series. Plainly these are not suitable for grade-school, but have the great benefit of making the calculus relationship $\frac{d}{d\theta}\sin\theta = \cos\theta$ very simple. Establishing this using the triangle definition is a somewhat tricky!

Measuring Angles

There are two standard ways to measure angles (to sensibly associate a *number* to each angle).

Degrees A full revolution has 360° and a right-angle 90°. Degree measure dates back to ancient Babylon 2–4000 years ago.[7]

Radians The radian measure of an angle is the length of the arc subtending the angle in a circle of radius 1. Since the circumference of a unit circle is 2π, we have the following identifications.

Degrees	Radians	$\sin\theta$	$\cos\theta$	$\tan\theta$
0°	0	0	1	0
30°	$\frac{\pi}{6}$	$\frac{1}{2}$	$\frac{\sqrt{3}}{2}$	$\frac{1}{\sqrt{3}}$
45°	$\frac{\pi}{4}$	$\frac{1}{\sqrt{2}}$	$\frac{1}{\sqrt{2}}$	1
60°	$\frac{\pi}{3}$	$\frac{\sqrt{3}}{2}$	$\frac{1}{2}$	$\sqrt{3}$
90°	$\frac{\pi}{2}$	1	0	n/a
180°	π	0	-1	0

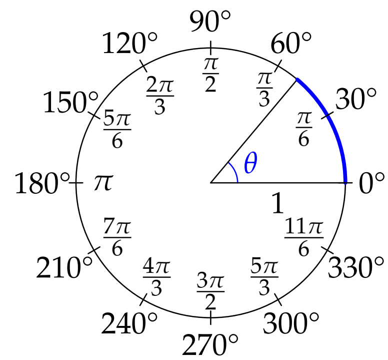

In elementary mathematics, degrees are the most common way to measure angles. Do you know any other methods?

Exercises 2.1. 1. The identity $\cos^2\theta + \sin^2\theta = 1$ is the Pythagorean Theorem in disguise. Why?

2. The word *sine* is the result of a long list of translations and transliterations from an ancient Sanskrit term meaning *half-chord*. For the chord picture on page 21, how does the length of the chord $\operatorname{crd}\theta$ relate to modern trigonometric functions?

3. It is conventional not to state units when using radians since they are effectively a ratio and therefore *unitless*. Think this through: if the central angle in a circle of radius r is subtended by an arc with arc-length ℓ, what is the radian measure of the angle? What facts from Euclidean geometry justify this observation?

4. Explain how to get the values of sine and cosine in the above table.
(*Hint: Draw some triangles and use Pythagoras*!)

5. Using the pictures, explain why we have the relations

$$\sin(\tfrac{\pi}{2} - \theta) = \cos\theta = \sin(\theta + \tfrac{\pi}{2}), \qquad \sin(-\theta) = -\sin\theta, \qquad \sin(\pi - \theta) = \sin\theta$$

(*You cannot use multiple-angle formulæ for this!*)

[7] It is not known why they chose 360, but it fits nicely with their *base-60* system of counting (decimals are base-10). The traditional subdivisions of a degree are also base-60. For instance, 34°12′45″ is 34 degrees, 12 (arc)minutes and 45 (arc)seconds; converted to decimal notation, this becomes

$$34°12'45'' = 34 + \frac{12}{60} + \frac{45}{60^2} = 34.2125°$$

The standard hour-minute-second measurement of time has the same origin.

2.2 Periodicity, Graphs & Inverses

One advantage of the circle definition is that it makes sketching the graphs of sine and cosine very easy. Simply draw axes next to a unit circle and transfer the heights across.

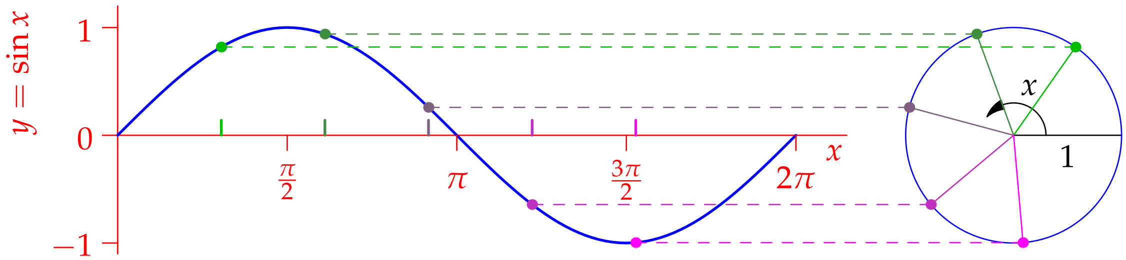

By Exercise 2.1.5, the graph of $\cos x = \sin(x + \frac{\pi}{2})$ is simply that of sine shifted $\frac{\pi}{2} = 90°$ to the left. Moreover, the circle definition allows us easily to extend trigonometric functions *periodically* since we can measure the polar angle by looping as many times round the origin as we like: for any integer n,

$$\sin(\theta + 2n\pi) = \sin\theta, \quad \cos(\theta + 2n\pi) = \cos\theta$$

Otherwise said, sine and cosine have period 2π radians (360°).

Sine and cosine are *non-invertible* unless we choose a domain on which they are 1–1.

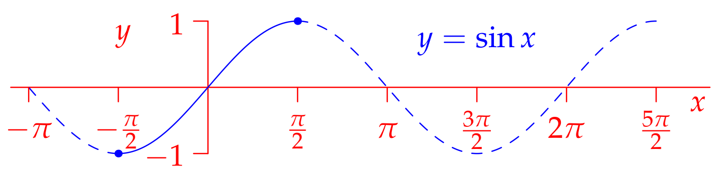

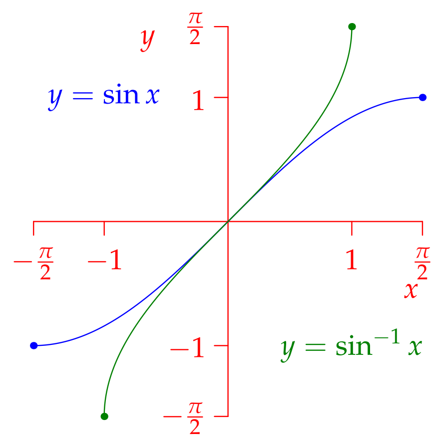

$f(x) = \sin x$ is 1–1 on the domain $[-\frac{\pi}{2}, \frac{\pi}{2}]$

Inverse function $f^{-1}(x) = \arcsin x = \sin^{-1} x$

Domain $\text{dom}(\arcsin) = [-1, 1] = \text{range}(\sin)$

Range $\text{range}(\arcsin) = [-\frac{\pi}{2}, \frac{\pi}{2}] = \text{dom}(\sin)$

This is why your calculator always returns a value in the interval $[-\frac{\pi}{2}, \frac{\pi}{2}] = [-90°, 90°]$ when you hit the $\sin^{-1}$ button.

Example 2.2. If you know the graphs, then symmetry and periodicity help you solve equations. For example, if $\sin\theta = \frac{9}{10}$ then all solutions are given by

$$\theta = \sin^{-1}\frac{9}{10} + 2\pi n \quad \text{or} \quad \pi - \sin^{-1}\frac{9}{10} + 2\pi n \qquad (n \text{ is any integer})$$

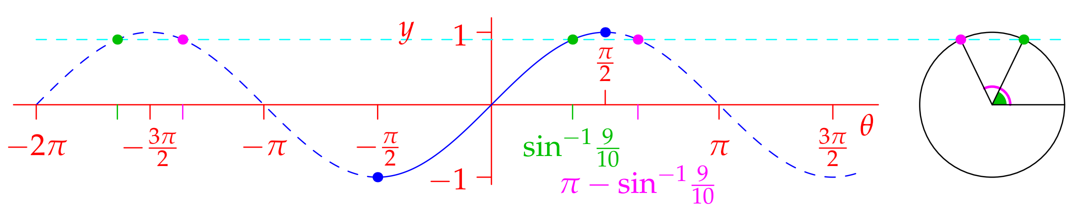

Alternatively, we could use the circle definition directly: $\sin\theta = \frac{9}{10}$ means we want angles θ corresponding to the intersections of the unit circle with the *horizontal line* $y = \frac{9}{10}$.

Periodic Models Trig functions find applications in modeling precisely because they are *periodic*. In general, a function has period T if

$$f(x+T) = f(x) \quad \text{for all } x$$

It is easy to find the period of the function $f(x) = \sin kx$ just by considering what we have to add to the input x to increase the argument kx of sine by 2π:

$$T = \frac{2\pi}{k} \implies f(x+T) = \sin(kx+2\pi) = \sin kx = f(x)$$

We may therefore obtain a simple periodic model regardless of what period is required.

Example 2.3. On a given day, high tide occurs at 2:00 with a water depth of 10 ft, whereas low tide occurs at 8:12 with a depth of 4 ft. We might model this using a periodic function with period $T = 2 \times 6\frac{12}{60} = \frac{62}{5}$ hours. For instance

$$h(t) = 7 + 3\cos\left(\frac{5\pi}{31}(t-4)\right)$$

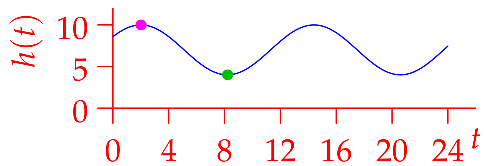

where t is measured in hours from midnight might be suitable. In reality, tidal height is very close to being periodic, but the magnitude of the high and low tides are somewhat variable.

In fact *any* periodic function may be approximated using trigonometric functions. Indeed if $f(x)$ has period T and we define constants

$$a_n = \int_{-\frac{T}{2}}^{\frac{T}{2}} f(x)\cos\frac{2\pi nx}{T}\,dx, \qquad b_n = \int_{-\frac{T}{2}}^{\frac{T}{2}} f(x)\sin\frac{2\pi nx}{T}\,dx \tag{$*$}$$

then

$$f(x) \approx \frac{a_0}{2} + a_1\cos\frac{2\pi x}{T} + b_1\sin\frac{2\pi x}{T} + a_2\cos\frac{4\pi x}{T} + b_2\sin\frac{4\pi x}{T} + \cdots \tag{$\dagger$}$$

This is the *Fourier series* of $f(x)$. It often takes only a small number of terms to obtain a very good approximation. Modern data-compression algorithms often employ Fourier series. Given a periodic function $f(x)$, one uses a computer to estimate (say) the first 100 Fourier coefficients ($*$) and transmits these values to the receiver, who recovers an approximation to the original function using ($\dagger$).

Example 2.4. A square-wave function with period $T = 2\pi$ is given by

$$f(x) = \begin{cases} 1 & \text{if } 0 \le x < \pi \\ -1 & \text{if } \pi \le x < 2\pi \end{cases}$$

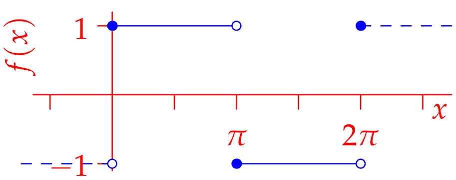

extended periodically to the real line. With a little calculus, it is easily checked that the Fourier coefficients are

$$a_n = \int_{-\pi}^{\pi} f(x)\cos nx\,dx = 0, \qquad b_n = \int_{-\pi}^{\pi} f(x)\sin nx\,dx = \begin{cases} \frac{4}{\pi n} & \text{if } n \text{ is odd} \\ 0 & \text{if } n \text{ is even} \end{cases}$$

Use a graphics tool to see how the first few terms of the series approximate the function.

Exercises 2.2. 1. $f(x) = \sin x$ is also 1–1 on the interval $[\frac{\pi}{2}, \frac{3\pi}{2}]$. Sketch the graph of its corresponding inverse function.

2. Draw the graph for cosine and observe that it is invertible if we restrict the domain to the interval $[0, \pi]$. Draw the graph of $\cos^{-1}$.

3. Describe all solutions to the equation $\cos x = -0.2$.

4. Explain why the tangent function has period π; that is $\tan(\theta + n\pi) = \tan\theta$. What facts are we using about sine and cosine and why are they obvious from the definition?

5. Describe all solutions to the equation $\tan x = 5$.

6. (a) Suppose $\theta = \cos^{-1}\frac{9}{41}$. Find the exact values for $\sin\theta$ and $\tan\theta$.
   (b) What changes if $\theta = \cos^{-1}\frac{-9}{41}$?

7. Let $f(x) = \csc x = \frac{1}{\sin x}$ be the cosecant function. Describe a domain on which this function is 1–1 and sketch the graph of its inverse $y = f^{-1}(x)$.

8. Use a computer to sketch the curve

$$y = 2\left(\sin x - \frac{1}{2}\sin 2x + \frac{1}{3}\sin 3x - \frac{1}{4}\sin 4x + \frac{1}{5}\sin 5x\right)$$

What simple periodic function do you think this is approximating?

2.3 Solving Triangles

Basic trigonometry often involves finding the edges and angles of a triangle given partial data.

Example 2.5. To find the height h of a tall tree, two angles of elevation 45° and 30° are measured a distance 20 ft apart along a straight line from the base of the trunk.

This is easily attacked by drawing a picture and observing that we have two right-triangles. If the (unknown) distance from the base of the tree to the nearer measurement is x, then

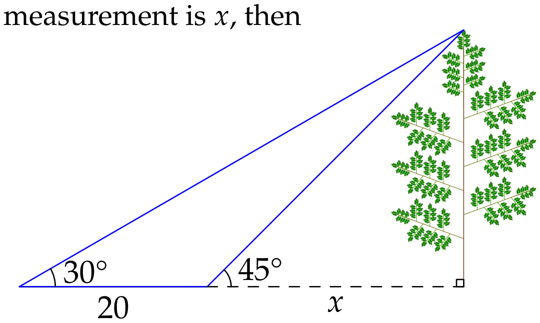

$$\frac{1}{\sqrt{3}} = \tan 30° = \frac{h}{x+20} \qquad 1 = \tan 45° = \frac{h}{x}$$

Substituting the second equation into the first returns

$$h = \frac{20}{\sqrt{3}-1} \approx 27.32\,\text{ft}$$

In fact there is enough data in the problem to recover everything about the original triangle.

- The second base angle is $(180° - 45° = 135°)$.
- The third (summit) angle is $180° - 30° - 135° = 15°$.
- Two applications of Pythagoras compute the remaining sides of the triangle

$$\sqrt{x^2+h^2} = \sqrt{2}h = \frac{20\sqrt{2}}{\sqrt{3}-1} \approx 38.64$$
$$\sqrt{h^2+(x+20)^2} = \sqrt{h^2+3h^2} = 2h = \frac{40}{\sqrt{3}-1} \approx 54.64$$

The example is just a disguised version of *solving a triangle*: computing all six sides and angles of a triangle given three of them. The Euclidean triangle congruence theorems tell us which combinations are sufficient to determine all the others. The example is the ASA congruence: angle-side-angle data (30°–20–135°) is enough to compute everything else about the triangle.

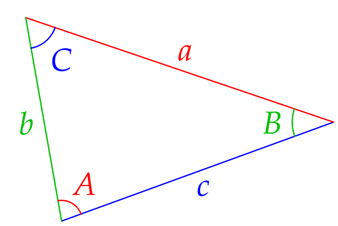

When in doubt, you can always attack basic trigonometry problems as we did in the example: create a right-triangle, then use the definitions of sin/cos/tan and/or Pythagoras.

Example 2.6. Given the SAS (side-angle-side) combination 5–60°–9, find the third side of the triangle.

The altitude h creates two right-triangles, from which

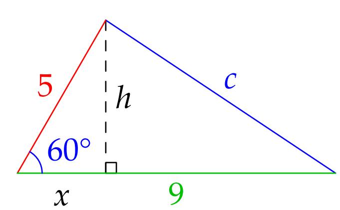

$$\begin{aligned} & h = 5\sin 60°, \quad x = 5\cos 60° \\ \implies & c^2 = (9-x)^2 + h^2 = 9^2 + (x^2+h^2) - 18x \\ & \quad = 9^2 + 5^2 - 18 \cdot 5\cos 60° = 61 \\ \implies & c = \sqrt{61} \approx 7.81 \end{aligned}$$

Since we now know c and $9 - x = \frac{13}{2}$ the remaining angles could also be easily found.

In elementary situations it is typically easier to have students drop the perpendicular as we've done. However, once comfortable with the method, it is helpful to have short-cuts which skip the need to work with the perpendicular at all.

Theorem 2.7 (Sine and Cosine Rules). *For any triangle,*

$$\frac{\sin A}{a} = \frac{\sin B}{b} = \frac{\sin C}{c} \quad \text{and} \quad c^2 = a^2 + b^2 - 2ab\cos C$$

The cosine rule is just the Pythagorean Theorem with a correction for non-right triangles. Both rules follow straightforwardly by drawing an altitude as before!

Proof. Consider the picture. We have

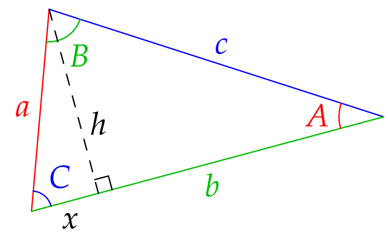

$$h = a\sin C = c\sin A, \quad x = a\cos C, \quad b - x = c\cos A$$

The first equation rearranges to

$$\frac{\sin A}{a} = \frac{\sin C}{c}$$

Two applications of Pythagoras give the cosine rule

$$c^2 = h^2 + (b-x)^2 = h^2 + x^2 + b^2 - 2bx = a^2 + b^2 - 2ab\cos C$$

The remaining part of the sine rule and the other versions of the cosine rule are obtained by choosing other altitudes. ∎

Here are two examples where we use the rules instead of explicitly drawing an altitude.

Examples 2.8. 1. A triangle has sides 2 and $\sqrt{3}-1$, and the angle between them is 120°. Find the remaining sides and angles.

We apply the cosine rule with $a = 2$, $b = \sqrt{3} - 1$ and $C = 120°$

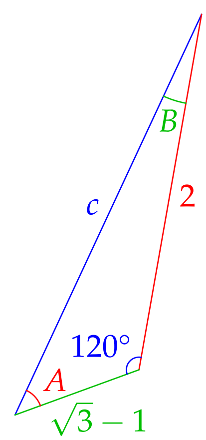

$$\begin{aligned} c^2 &= a^2 + b^2 - 2ab\cos C \\ &= 2^2 + (\sqrt{3}-1)^2 - 2\cdot 2(\sqrt{3}-1)\cos 120° \\ &= 4 + 3 + 1 - 2\sqrt{3} + 2(\sqrt{3}-1) = 6 \end{aligned}$$

We have an opposite pair $(c, C) = (\sqrt{6}, 120°)$, so the sine rule may be used

$$\sin A = \frac{2}{\sqrt{6}}\sin 120° = \frac{2\sqrt{3}}{2\sqrt{6}} = \frac{1}{\sqrt{2}} \implies A = 45°$$

We chose the acute angle since $A = 180° - B - C = 60° - B < 90°$.
The final angle is then $B = 180° - 45° - 120° = 15°$.

You could instead drop a perpendicular, say from the vertex A to the *extension* of the side of length 2. Think about why the perpendicular has to be *outside* the triangle…

2. A triangle has one side with length 5 and its two adjacent angles are 40° and 65°. Find the remaining data.

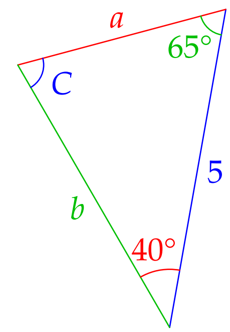

This time the initial data is ASA. Writing $c = 5$, $A = 40°$ and $B = 65°$, the remaining angle is plainly

$$C = 180° - 40° - 65° = 75°$$

This gives us an opposite pair (c, C), so we can apply the sine rule

$$a = c\frac{\sin A}{\sin C} = 5\frac{\sin 40°}{\sin 75°} \approx 3.327$$

A second application yields

$$b = c\frac{\sin B}{\sin c} = 5\frac{\sin 65°}{\sin 75°} \approx 4.691$$

3. (Discuss: courtesy of an 8 year-old contributor) Model Earth as a sphere of radius 3963 miles. If identical vertical ladders are placed in Irvine, CA and Irvine, Scotland, 5145 miles apart by great circle, how tall would they have to be for people at the top to 'see' each other?

Multiple-angle Formulae

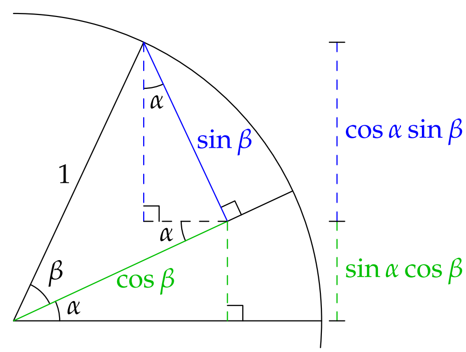

Also useful in the context of basic trigonometry is the ability to sum angles. The picture provides a simple justification of

$$\sin(\alpha + \beta) = \sin\alpha\cos\beta + \cos\alpha\sin\beta$$

at least when $0 < \alpha + \beta < \frac{\pi}{2}$. If you look carefully, you should be able to see how the same picture establishes

$$\cos(\alpha + \beta) = \cos\alpha\cos\beta - \sin\alpha\sin\beta$$

Exercises 2.3. 1. Find the remaining angles in the triangle in Example 2.6.

2. The other Euclidean congruence theorems are SSS and SAA. Explain how to solve triangles using these minimal data in two ways:

   (a) By drawing an altitude. (b) Using the sine/cosine rules.

3. SSA isn't a triangle congruence theorem. For instance, there are *two* non-congruent triangles with data $a = 1$, $b = \sqrt{3}$ and $A = 30°$. Find them.

4. Use the multiple-angle formulae to derive the familiar expressions for $\sin 2\theta$ and $\cos 2\theta$.

5. Find the exact value of $\sin 105°$.

6. (a) Find an expression for $\tan(\alpha + \beta)$ purely in terms of $\tan\alpha$ and $\tan\beta$.

   (b) Two wooden wedges with slope $\frac{1}{4}$ are placed on top of each other to make a steeper slope. What is the gradient of the new slope?

7. Carefully explain why the answer to Example 2.8.3 is approximately 1012 miles.

2.4 Polar Co-ordinates

Definition 2.1 provides an alternative way to describe points in the plane. If θ is the polar angle of a point with Cartesian (rectangular) co-ordinates (x, y), then its polar-coordinates are precisely the values (r, θ) seen in the definition!
Computing $x = r\cos\theta$ and $y = r\sin\theta$ is easy given r and θ.

Example 2.9. A point with polar co-ordinates $(r, \theta) = (2, \frac{5\pi}{6})$ has Cartesian co-ordinates

$$(x, y) = \left(2\cos\tfrac{5\pi}{6}, 2\sin\tfrac{5\pi}{6}\right) = (-\sqrt{3}, 1)$$

Computing polar co-ordinates from Cartesian is harder, requiring some *visualization.*

1. Every point (x, y) has a unique radius $r = \sqrt{x^2 + y^2}$, but not polar angle. If θ is a polar angle, so is $\theta + 2\pi n$ for any integer $n \in \mathbb{Z}$. The origin $(x, y) = (0, 0)$ is even stranger; certainly $r = 0$, but *any* θ is a legitimate polar angle!

2. Whenever $x \neq 0$ (away from the y-axis),

$$\begin{cases} x = r\cos\theta \\ y = r\sin\theta \end{cases} \implies \tan\theta = \frac{y}{x}$$

   however, this *doesn't* guarantee that $\theta = \tan^{-1}\frac{y}{x}$. Continuing the example shows us why...

Example (2.9, cont). If $(x, y) = (-\sqrt{3}, 1)$, then the radius is easy

$$r = \sqrt{(\sqrt{3})^2 + 1^2} = 2$$

For the polar angle,

$$\tan\theta = \frac{y}{x} = -\frac{1}{\sqrt{3}} = \tan\left(-\frac{\pi}{6}\right) \not\Longrightarrow \theta = -\frac{\pi}{6}$$

Arctan has range $(-\frac{\pi}{2}, \frac{\pi}{2})$, so always returns an angle in quadrants 1 or 4. Our point is in the *second* quadrant $(x < 0 < y)$ so we need to adjust, using the fact that tan is π-periodic:

$$\theta = \pi - \frac{\pi}{6} = \frac{5\pi}{6} = 150°$$

We could alternatively add any integer multiple of 2π.

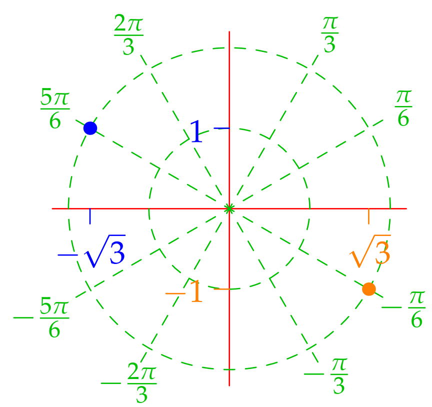

The example wasn't too tricky since the polar angle was exactly computable. When you have to rely on a calculator, it is much easier to make a mistake.

Example 2.10. The point $(x, y) = (-8, -15)$ has polar co-ordinates (quadrant 3!)

$$r = \sqrt{8^2 + 15^2} = 17, \qquad \theta = \pi + \tan^{-1}\tfrac{15}{8} \approx 241.93°$$

We could summarize with formulæ describing precisely how to compute θ dependent on quadrant (the signs of x, y), though it is better to get in to the habit of drawing a picture!

Curves in Polar Co-ordinates

Polar co-ordinates are well-suited to describing curves that encircle the origin. Indeed circles centered at the origin with radius $a > 0$ have the very simple polar form $r = a$. Converting to rectangular co-ordinates recovers the the natural parametrization of a circle:

$$x(\theta) = a\cos\theta, \quad y(\theta) = a\sin\theta$$

This partly explains why mathematicians call sine and cosine *circular functions*.

General polar graphs are harder to visualize, though the major reason is lack of familiarity. Have a little empathy: to graph *polar* functions, you'll likely have to follow the same approach as new students use to sketch *Cartesian* curves like $y = x^2$! Here are a couple of examples.

Examples 2.11. 1. The curve $r = \theta$ is relatively easy to plot since r increases at exactly the same rate as the angle; we therefore have a *spiral*.

To confirm this, plot several points (θ, θ); we've done for θ in multiples of $\frac{\pi}{6}$ (30°) from 0 to 2π. It is sensible to use 'polar graph paper' with concentric circles separated by (say) $\frac{\pi}{2} \approx 1.57$.

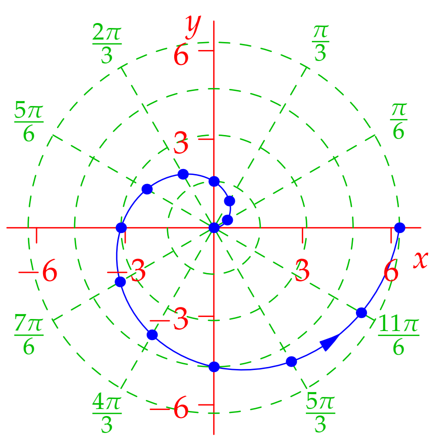

$r = \theta, \ 0 \le \theta \le 2\pi$

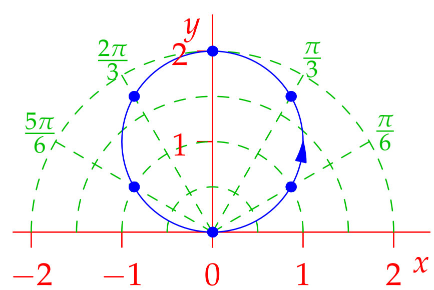

$r = 2\sin\theta, \ 0 \le \theta \le \pi$

θ	$\frac{\pi}{6}$	$\frac{2\pi}{3}$	$\frac{\pi}{2}$	$\frac{4\pi}{3}$	$\frac{5\pi}{6}$	π
$2\sin\theta$	1	1.73	1	1.73	1	0

2. The curve $r = 2\sin\theta$ is a little easier to work with since we know exact values for sine, assisted by $\sqrt{3} \approx 1.73$.

This looks like a circle! To see this, multiply both sides by r and complete the square:

$$r^2 = 2r\sin\theta \implies x^2 + y^2 = 2y \implies x^2 + (y-1)^2 = 1$$

describes the set of points with distance 1 from the point $(0, 1)$: a circle!

You should think about what happens in both examples if we extend the domain:

- What would $r = \theta$ look like if θ were allowed to be *negative*?
- What happens to $r = 2\sin\theta$ when $\theta > \pi$?

Exercises 2.4. 1. Convert the following points to polar co-ordinates.

(a) $(-5, 5)$ (b) $(3, -4)$ (c) $(-5\sqrt{3}, -15)$
(d) $(-1, \tan 3)$ (tricky—this is 3 radians!)

2. If $a > 0$, describe the curve with polar equation $r = 2a\cos\theta$.
(*Be careful with* $\theta > \frac{\pi}{2}$ *since cosine goes negative...*)

3. The algebraic trickery in the last example sometimes bears fruit, though you have to be lucky! By multiplying both sides by $1 - \sin\theta$ and converting to rectangular co-ordinates, show that the polar function

$$r(\theta) = \frac{2a}{1 - \sin\theta}$$

is a parabola in disguise and sketch it when $a = 1$. How does the graph depend on a?

4. In a similar vein to the previous question, sketch the curve $r = \frac{2}{2+\sin\theta}$. What type of curve is this?

5. Try to sketch the following curves.

(a) $r = \theta(\theta - 4)$ (b) $r = (\theta - 1)^2 + 1$ (c) $r = (\theta - 1)^2 - 1$

As well as plotting points directly, you should sketch the curve first on rectangular axes (e.g., (a) is $y = x(x - 4)$). What happens to (c) when $\theta = 1$?

Once you've tried these, use a grapher to see if you're right, though see how close you can get without it!

3 Exponential and Logarithmic Functions & Models

Introducing exponential functions without calculus presents a significant challenge. The simplest approach is as a short-hand notation for *repeated multiplication*: for instance

$$a^5 = a \cdot a \cdot a \cdot a \cdot a$$

analogous to how multiplication represents repeated addition

$$5a = a + a + a + a + a$$

The problem with this approach is that it doesn't help you understand what should be meant by, say, $a^{3/4}$ or $a^{\sqrt{2}}$: multiplying something by itself '$\sqrt{2}$ times' sounds[8] insane!

To rigorously address this problem requires *continuity* and other ideas surrounding the foundations of calculus which you'll encounter in upper-division analysis; topics unsuitable for this course. Instead, we assume some familiarity with exponential functions via introductory calculus, where they are unavoidable and offer two ways to introduce exponential functions and e via modelling.

3.1 The Natural Growth Model

A basic model for any variable quantity is that its *rate of change be proportional to the quantity itself.* This idea necessarily needs some calculus; as a differential equation,

$$\frac{\mathrm{d}y}{\mathrm{d}x} = ky$$

where k is a constant; if $k > 0$ this is the *natural growth equation*, if $k < 0$ the *natural decay equation.* This is commonly encountered when modelling *population growth*; an otherwise unconstrained population seems like its growth rate should be proportional to its size (twice the people, twice the babies...). This model is hugely applicable, since *population* can refer to essentially any quantifiable value: people, bacteria, money, reagents in a chemical/nuclear reaction, etc.

Example 3.1. The simplest natural growth equation has $k = 1$:

$$\frac{\mathrm{d}y}{\mathrm{d}x} = y$$

If a point (x, y) lies on a solution curve, then the differential equation tells us the *direction of travel* of the solution. We may visualize this by drawing an arrow with slope $\frac{\mathrm{d}y}{\mathrm{d}x} = y$; the arrows are *tangent* to any solution.[9] You should easily be able to sketch some other solution curves.

You should, of course, recognize the graph...

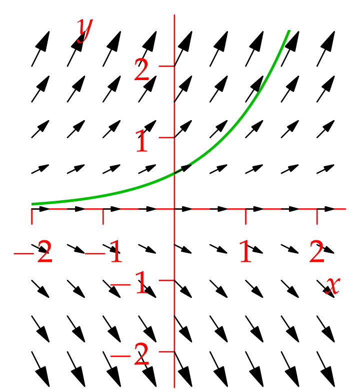

[8] The same issue arises for multiplication: $3\sqrt{2} = \sqrt{2} + \sqrt{2} + \sqrt{2}$ is relatively easy to understand, but how would you convince someone what $\pi\sqrt{2}$ means?

[9] A similar approach is available for any first-order differential equation $\frac{\mathrm{d}y}{\mathrm{d}x} = F(x, y)$: the equation defines its *slope field* (arrows), to which solution curves must be tangent.

Definition 3.2. Let $a > 0$ be constant. The *exponential function with base a* is $f(x) = a^x$.

Recall the *exponential laws*, which are very natural when x, y, r are positive integers:

$$a^{x+y} = a^x a^y \qquad a^{x-y} = \frac{a^x}{a^y} \qquad (a^x)^r = a^{rx}$$

These hold for all exponents, with the same continuity caveats we saw previously. For modelling, the crucial property of exponential functions is that they have proportional derivative.

Theorem 3.3. *The rate of change of $f(x) = a^x$ is proportional to $f(x)$. Specifically,*

$$f'(x) = \lim_{h\to 0} \frac{a^{x+h} - a^x}{h} = a^x \lim_{h\to 0} \frac{a^h - 1}{h}$$

so that $f(x) = a^x$ satisfies the natural growth/decay equation $\frac{dy}{dx} = ky$ with proportionality constant

$$k = f'(0) = \lim_{h\to 0} \frac{a^h - 1}{h}$$

Example 3.4. We estimate the proportionality constant $k = \lim_{h\to 0} \frac{a^h - 1}{h}$ to 3 d.p. using a calculator for four values of h:

a	2	2.5	2.7	2.75	3	5
$\frac{a^{0.1}-1}{0.1}$	0.718	0.960	1.044	1.065	1.161	1.746
$\frac{a^{0.01}-1}{0.01}$	0.696	0.921	0.998	1.017	1.105	1.622
$\frac{a^{0.001}-1}{0.001}$	0.693	0.917	0.994	1.012	1.099	1.611
$\frac{a^{0.0001}-1}{0.0001}$	0.693	0.916	0.993	1.012	1.099	1.610

What is happening to the proportionality constant as a increases?

It appears as if there is a special number somewhere between 2.7 and 2.75 for which the proportionality constant is precisely $k = 1$.

Definition 3.5. The value $e = 2.71828\ldots$ is the unique real number such that $\lim_{h\to 0} \frac{e^h-1}{h} = 1$. The *natural*[10] *exponential function* $\exp(x) = e^x$ has derivative $\frac{d}{dx}e^x = e^x$.

[10] *Natural* here means *unavoidable*: an old cliche suggests that if aliens were to land on Earth, they'd have to understand e given the technology they'd require to get here. Of course they'd likely use a different symbol; ours comes from Leonhard Euler around 1728. Like π and $\sqrt{2}$, the constant e is an *irrational number*: its decimal representation contains no repeating pattern. There isn't the same geeky fascination with memorizing the digits of e as there is with π, neither is there an 'e-day' (Feb 7th at 6:28 p.m.?).

The function $f(x) = \frac{1}{2}e^x$ is plotted in Example 3.1. Of course there are many other solutions to the natural growth equation $\frac{dy}{dx} = y$: for any constants c, k,

$$y = ce^{kx} \implies \frac{dy}{dx} = \frac{d}{dx}ce^{kx} = cke^{kx} = ky$$

In fact the converse also holds; for the details, take a differential equations course!

Theorem 3.6. *The solutions to the natural growth equation $\frac{dy}{dx} = ky$ are precisely the functions*

$$y(x) = y_0 e^{kx} = y_0 \exp(kx)$$

where $y_0 = y(0)$ is the initial value.

Example 3.7. A Petri dish contains a population $P(t)$ of bacteria satisfying the natural growth equation $\frac{dP}{dt} = 0.5P$ where time is measured in weeks from the start of the year.

If $P(0) = 100$ bacteria, then $P(t) = 100e^{0.5t}$. Specifically, at the end of January ($4\frac{3}{7}$ weeks) one expects there to be

$$P(\tfrac{31}{7}) = 100 \exp \tfrac{31}{14} = 915 \text{ bacteria}$$

Note that the exponential doesn't return 915 exactly; this is only an approximation. Models like this work best for large populations where integer rounding errors are of minimal concern.

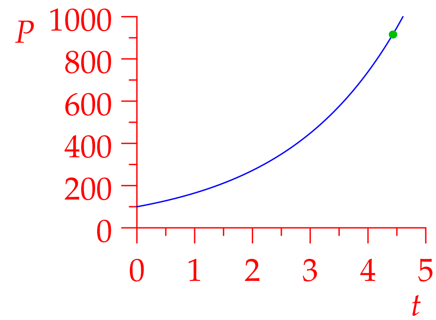

Compound Interest and the Discovery of e

The first description of e came in 1683 when Jacob Bernoulli tried to model the growth of money in a hypothetical bank account. We give a modernized version of his approach.

Example 3.8. \$1 is deposited in an account paying 100% interest per year (nice!). Bernoulli observed that the money in the account at the end of the year depends on *when* the interest is paid.

- If the interest is paid once at the end of the year (this is called *simple* interest), you'll have \$2.
- If half the interest (50¢) is paid at six months, then the balance (\$1.50) earns $\frac{1}{2} \cdot 1.50 = 75$¢ interest for the rest of the year; you'll finish the year with \$2.25 in the account.
- If the interest is paid in four installments, we have the following table of transactions (data is rounded to the nearest cent)

Date	Interest Paid	Balance
1st Jan	—	\$1
1st Apr	25¢	\$1.25
1st July	$\frac{1}{4} \cdot 1.25 = 31$¢	\$1.56
1st Oct	$\frac{1}{4} \cdot 1.56 = 39$¢	\$1.95
New Year	$\frac{1}{4} \cdot 1.95 = 49$¢	\$2.44

More succinctly, the year-end balance is $\left(1 + \frac{1}{4}\right)^4 = \$2.44$.

- More generally, if the interest is paid over n equally spaced intervals, the account balance at the end of the year would be $\$(1+\frac{1}{n})^n$. Here are a few examples rounded things to 5 d.p.

Frequency	Balance after 1 year ($)
Every month	$(1+\frac{1}{12})^{12} = 2.61304$
Every day	$(1+\frac{1}{365})^{365} = 2.71457$
Every hour	$(1+\frac{1}{8760})^{8760} = 2.71813$
Every second	$(1+\frac{1}{31536000})^{31536000} = 2.71828$

As the frequency of payment increases, it appears as if the balance is increasing to $\$e\ldots$

In fact this is a theorem, though it requires significant work (beyond this class) to prove it:

$$e = \lim_{n\to\infty}\left(1+\frac{1}{n}\right)^n \quad \text{and more generally} \quad e^x = \lim_{n\to\infty}\left(1+\frac{x}{n}\right)^n$$

This again shows that e arises very naturally.

Simple, Monthly & Continuous Interest In finance, interest is typically computed in one of three ways. In each case we describe the result of investing $1 at an annual interest rate of $r\% = \frac{r}{100}$.

Simple interest You are paid $\frac{r}{100}$ dollars at the end of the year. Your invested dollar becomes $1+\frac{r}{100}$ dollars.

Monthly interest Each month you are paid $\frac{r}{12}\%$ of your current balance. This amounts to a balance of $(1+\frac{r}{1200})^{12}$ dollars at year's end. The period need not be monthly: if interest is paid in n installments, the balance would be $(1+\frac{r}{100n})^n$.

Continuous interest After t years (can be any fraction of a year!) your dollar-balance is

$$e^{\frac{rt}{100}} = \exp\frac{rt}{100} = \lim_{n\to\infty}\left(1+\frac{rt}{100n}\right)^n$$

Example 3.9. A bank account earns 6% annual interest paid monthly. To what simple annual interest rate does this correspond? Would you perfer an account paying 6% continuously?

At the end of the year, $1 becomes

$$(1+0.0612)^{12} = 1.005^{12} \approx 1.06168\ldots$$

corresponding to a simple interest rate of 6.17%. By cobntrast, 6% continuous interest would result in your dollar becoming $e^{6/100} \approx 1.06184$, corresponding to a (marginally) higher simple interest rate of 6.18%. You should prefer this, particularly if you have a lot of money to invest! The difference is more noticable with an investment of $1000 over ten years:

$$1000 \times 1.005^{120} = \$1819.40 \quad \text{versus} \quad 1000e^{0.6} = \$1822.12$$

There are several reasons for these varying approaches, not all of them consumer-friendly:

1. Simple interest is simple! It is easy to understand and compute, but hard to decide how or even whether to compute interest for parts of a year.
2. Monthly interest fits with most paychecks, so is sensible for loans, particularly mortgages.
3. Continuous interest allows the balance of an account to be found easily at any time, even between interest payment dates. It is also much easier to apply mathematical analysis (calculus).
4. A company can make an interest rate appear *higher* (if a savings account) or *lower* (if a loan) by choosing which way to quote an interest rate.

Example 3.10. A bank quotes you a loan with a continuously compounded interest rate of 7%. If you borrow \$100,000, then at the end of the year you'll owe

$$100000e^{0.07} = \$107,250.82$$

not the \$107,000 you might have expected! This corresponds to a simple interest rate (one payment at the end of the year) of 7.25%.[11]

Exercises 3.1. 1. Draw a slope field for the natural decay equation $\frac{dy}{dx} = -\frac{1}{3}y$ and use it to sketch the solution curve with initial condition $y(0) = 6$. What is the *function* $y(x)$ in this case?

2. Which of the following would you prefer for a savings account? Why?
   - 5% interest paid continuously.
   - 5.05% compounded monthly.
   - 5.1% paid at the end of the year.

3. You invest \$1000 in an account that pays 4% simple interest per year.
   (a) How much money will you have after 5 years?
   (b) If you close the account after 2 years and 3 months, the bank needs to decide how much interest to credit you with. Do this in two ways (the answers will be different!):
      i. Compute using the simple interest rate for 2.25 years ($(1 + \frac{r}{100})^{2.25}$).
      ii. Suppose that interest is paid at 4% for all completed years and then at 4% paid monthly for any completed months of an incomplete year. Find the balance of the account at closing.

4. Explain why the proportionality constant for $\left(\frac{1}{a}\right)^x$ is *negative* that for a^x: that is,

$$\lim_{h\to 0} \frac{(\frac{1}{a})^h - 1}{h} = -\lim_{h\to 0} \frac{a^h - 1}{h}$$

   Try to find both an *algebraic* reason and a *pictorial* one.

5. Sketch the function $f(x) = e^{-x^2}$. Where have you seen this before and what uses does this function have?

[11] In the US, mortgage companies typically quote an interest rate which they use to compound *monthly*. For example, if the quoted rate is 7%, then the effective annual (simple) interest rate is $\left(1 + \frac{0.07}{12}\right)^{12} - 1 = 7.229\%$. By law, this higher *effective APR* must be quoted somewhere, though it is unlikely to be as prominently posted...

3.2 Logarithmic Functions

Since $e > 1 > 0$, the natural exponential function satisfies several properties:

$$\lim_{x\to-\infty} e^x = 0, \qquad \lim_{x\to\infty} e^x = \infty, \qquad \frac{d}{dx}e^x = e^x > 0$$

Thus $\exp : \mathbb{R} \to (0,\infty)$ is a *differentiable (so continuous), increasing* function with domain $\mathbb{R}$ and range $(0,\infty)$. It is therefore *invertible*.

Definition 3.11. The *natural logarithm* $\ln : (0,\infty) \to \mathbb{R}$ is the inverse function to the natural exponential. That is,

- If $x > 0$, then $e^{\ln x} = x$;
- If $y \in \mathbb{R}$, then $\ln e^y = y$.

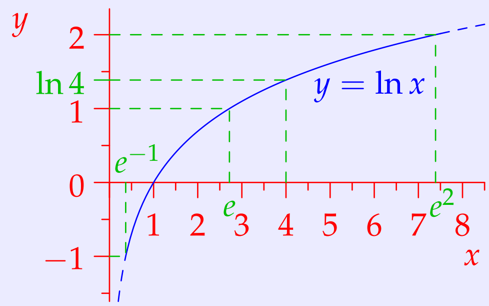

Since exp and ln are inverse functions, we can solve equations in the usual way: for instance,

$$e^{3x+1} = 100 \implies 3x+1 = \ln 100 \implies x = \frac{1}{3}(\ln 100 - 1) \approx 1.202$$

One of the great advantages of logarithms is that they allow every exponential function to be expressed in terms of the natural exponential: by the exponential laws,

$$a^x = (e^{\ln a})^x = e^{x\ln a}$$

This identity is crucial for interpreting and analyzing natural growth models.

Example 3.12. A population of rabbits doubles in size every 6 months. If there are 10 rabbits at the start of the year, how many rabbits do we expect there to be after 9 months, and how rapidly is the population increasing (births/month).

We are told that the population of rabbits after t months is

$$P(t) = 10 \cdot 2^{t/6}$$

After 9 months the population will be approximately

$$P(9) = 10 \cdot 2^{3/2} = 20\sqrt{2} \approx 28.28 \text{ rabbits}$$

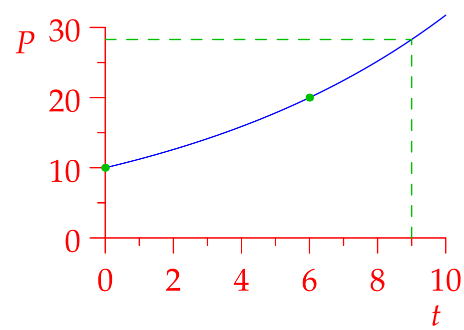

Moreover,

$$\frac{d}{dt}P(t) = \frac{d}{dt}10e^{\frac{t}{6}\ln 2} = \frac{10\ln 2}{6}2^{t/6} \implies P'(9) = \frac{10\ln 2}{6}2^{3/2} \approx 3.27 \text{ rabbits/month}$$

If you ask students this question, what do you expect to be the most common *incorrect* answers? Why?

The Logarithm Laws and General Logarithms

The logarithm laws should be familiar; they follow immediately from the above definition and the exponential laws (page 33)

$$e^{\ln x+\ln y} = e^{\ln x}e^{\ln y} = xy = e^{\ln xy} \implies \boxed{\ln xy = \ln x + \ln y}$$

$$\text{Similarly} \quad \boxed{\ln \tfrac{x}{y} = \ln x - \ln y \quad \text{and} \quad \ln x^r = r\ln x} \qquad (*)$$

These laws allow us to solve more general exponential equations.

$$2^x = 5 \implies x\ln 2 = \ln 5 \implies x = \frac{\ln 5}{\ln 2} \approx 2.322$$

More generally, if $a > 0$ and $a \neq 1$, then the exponential function with base a is invertible:

$$y = f(x) = a^x = e^{x\ln a} \implies \ln y = x\ln a \implies x = \frac{\ln y}{\ln a} \implies f^{-1}(x) = \frac{\ln x}{\ln a}$$

Definition 3.13. Given $a > 0$ and $a \neq 1$, the *logarithm with base a* is the function

$$\log_a x := \frac{\ln x}{\ln a}$$

As the inverse of the base a exponential function $y = a^x$, the base a logarithm satisfies

- If $x > 0$, then $a^{\log_a x} = x$;
- If $y \in \mathbb{R}$, then $\log_a a^y = y$.

The natural logarithm has base e. Unless the base is very simple (e.g. $a = 2$ or 10), we typically stick to using the natural logarithm. On a calculator, the 'log' button means $\log_{10}$.

Exercises 3.2.

1. Find the solution to the equation $4^{2-\sqrt{x}} = 10$.
2. Find the value of x which satisfies the equation $4^{6x} = 8$. Your answer should not contain any logarithms...
3. Over one year, find the continuous interest rate $s\%$ corresponding to a simple rate of 5%.
4. $y = a^x$ satisfies the natural growth equation $\frac{dy}{dx} = ky$; what is the value of k?
5. Verify the remaining logarithm laws $(*)$.
6. By differentiating the expression $e^{\ln x} = x$, verify that $\frac{d}{dx}\ln x = \frac{1}{x}$.
7. Sketch a graph of the functions $f(x) = \log_2 x$ and $g(x) = \log_{0.5} x$. How are they related? What happens to the graph of $\log_a x$ as a changes?
8. Logarithms were originally invented not for calculus but to simplify and multiply large numbers. In the pre-calculator era, it was common for students to carry a book of *log tables* for this purpose. Look up a log table and investigate how to use it.

3.3 Modifying the Natural Growth Model

In this section we discuss several examples of exponential models motivated by real-world situations. Remember that *modelling* always involves some guesswork and assumptions, which necessarily come with trade-offs: simpler assumptions/models are easier to analyze, but tend to be less accurate. Modelling is always a part of a feedback loop:

> Data/theory suggest a model whose predictions are tested against real-world data, suggesting changes/improvements to the model.

Applied mathematicians typically desire a 'Goldilocks' model: complicated enough to make accurate predictions without being too complicated to use.

Newton's Law of Cooling

A just-poured cup of coffee at 210°F is left outside when the air temperature is 50°F. It seems obvious that the coffee will cool down slowly *towards* 50°F; but how?

To help decide how to model this, ask yourself some questions:

1. When should the rate of cooling be most rapid?
2. What happens to the rate of change in the long run (large time)?
3. Can you suggest a family of functions which behave in this manner?

Hopefully it seems reasonable to model this with a shifted exponential function, where the temperature $T(t)$ of the coffee at time t satisfies

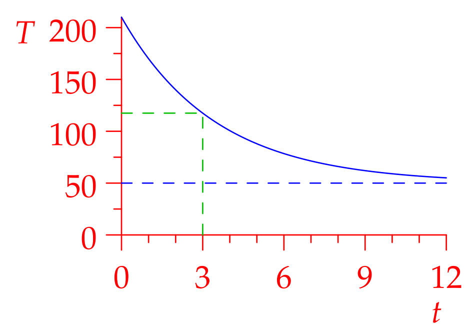

$$T(t) = 50 + 160e^{-kt}$$

for some positive constant k. This satisfy all our criteria:

- $T(0) = 210$°F.
- As t increases, e^{-kt} decreases to zero, so $T(t)$ decreases towards 50°F.
- The rate of cooling $\left|\frac{\mathrm{d}T}{\mathrm{d}t}\right| = 160ke^{-kt}$ is largest at $t = 0$ and decreases as t increases.

To complete the model, it is enough to supply one further data point.

Suppose after 2 minutes that the temperature of the coffee is 140°F. How long does it take for the coffee to cool to 100°F?

We know that $140 = T(2) = 50 + 160e^{-2k}$, whence

$$e^{-2k} = \frac{140 - 50}{160} = \frac{9}{16} \implies e^{-k} = \frac{3}{4} \implies T(t) = 50 + 160\left(\frac{3}{4}\right)^t$$

When $T(t) = 100$, we see that

$$\left(\frac{3}{4}\right)^t = \frac{100 - 50}{160} = \frac{5}{16} \implies t = \frac{\ln\frac{5}{16}}{\ln\frac{3}{4}} = \frac{\ln 16 - \ln 5}{\ln 4 - \ln 3} \approx 4.043 \text{ minutes}$$

This is an example of a general model called *Newton's law of cooling*, which asserts that the rate of temperature change of a body is proportional to the *difference* between the body and its surroundings. We have a simple modification of Theorem 3.6

Corollary 3.14. *If M and k are constant, then*

$$\frac{dy}{dt} = k(M - y) \iff y(t) = M + (y_0 - M)e^{-kt}$$

where $y_0 = y(0)$ is the initial value.

The Logistic Model

The natural growth model has one enormous drawback when applied to real-world populations: if $k > 0$, then a function $y(t) = y_0 e^{kt}$ *grows unboundedly*! In typical situations, environmental limitations (availability of food, water, space) mean that populations do not *explode* like this. The *logistic model* attempts to describe this phenomenon; it is based on two assumptions:

- When a population y is small, we want it to grow naturally $\frac{dy}{dt} \propto y$.
- We want y to approach a positive value M as $t \to \infty$.

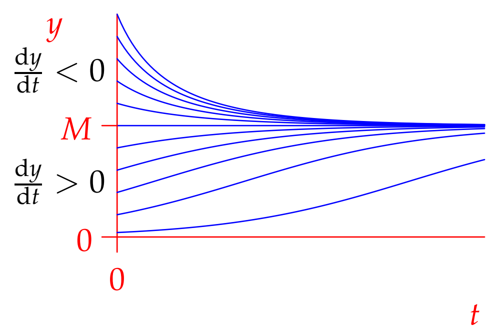

Given constants $k, M > 0$, the *logistic differential equation*

$$\frac{dy}{dt} = ky(M - y)$$

accomplishes both requirements. M is often referred to as the *carrying capacity* of the environment.

Theorem 3.15. *If $y_0 = y(0)$, then the solution to the logistic differential equation is*

$$y(t) = \frac{y_0 M}{y_0 + (M - y_0)e^{-kMt}}$$

You can check directly that this satisfies the differential equation just by differentiating, though it's a little ugly. If you've studied differential equations the method of separation of variables supplies the converse.

Example 3.16. A brewer pitches 100 billion yeast cells into a starter wort with the goal of growing it to 200 billion cells. After one hour, the wort contains 110 billion cells.

(a) How long must the brewer wait if we use a natural growth model?

(b) How long must the brewer wait if we use a logistic model where we also assume that the wort contains enough sugar to grow 250 billion yeast cells?

Let $P(t)$ be the yeast population in billions at time t hours. We therefore have $P(0) = 100$ and $P(1) = 110$, and want to find t such that $P(t) = 200$.

(a) The model is $\frac{dP}{dt} = ky$, which has solution

$$P(t) = P_0 e^{kt} = 100e^{kt}$$

Evaluating at $t = 1$ yields $1.1 = e^k$ whence

$$P(t) = 100(1.1)^t \implies t = \frac{\ln(0.01P)}{\ln 1.1} \approx 7.27 \text{ hours}$$

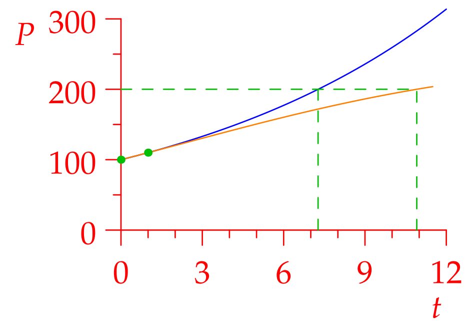

(b) The model is $\frac{dP}{dt} = ky(250 - y)$, with solution

$$P(t) = \frac{25000}{100 + (250 - 100)e^{-250kt}} = \frac{500}{2 + 3e^{-250kt}}$$

Evaluating at $t = 1$ yields

$$110 = \frac{500}{2 + 3e^{-250k}} \implies e^{250k} = \frac{33}{28}$$

whence

$$P(t) = \frac{500}{2 + 3\left(\frac{28}{33}\right)^t} \implies t = \frac{\ln 6}{\ln \frac{33}{28}} \approx 10.91 \text{ hours}$$

The logistic model is easily generalized.

Example 3.17. The population $P(t)$ (in 1000s) of fish in a lake obeys the logistic equation

$$\frac{dP}{dt} = \frac{1}{16}P(10 - P)$$

where t is measured in months. The first graph shows how the population recovers over a year if it starts at 2500 fish.

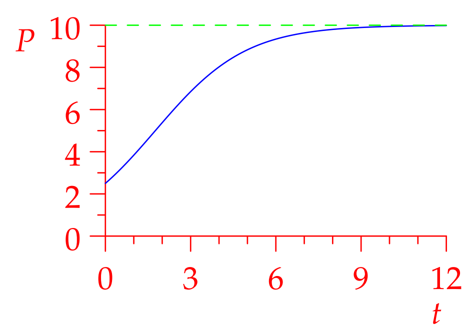

Now suppose 1000 fish are 'harvested' from the lake each month. The new model is then

$$\frac{dP}{dt} = \frac{1}{16}P(10 - P) - 1 = -\frac{1}{16}(P^2 - 10P + 16)$$
$$= -\frac{1}{16}(P - 2)(P - 8)$$

Substituting $Q = P - 2$, this is again logistic!

$$\frac{dQ}{dt} = \frac{1}{16}Q(6 - Q)$$

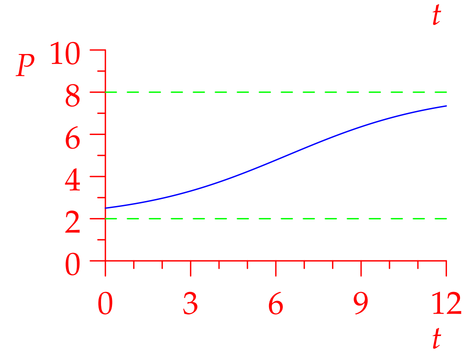

Provided the initial population $P(0) = Q(0) + 2$ is greater than 2000 fish, we expect the population to eventually stabilize at 8000 fish, though it takes a long time to get close to this if we start, as in the second graph, with only a little over 2000 fish.

Public-health Interventions

A population of 10,000 people is exposed to a novel virus. The best scientific understanding is that 1% of the susceptible population per day contracts the virus, the effects of the illness last ten days, after which a patient recovers and is immune from reinfection.

1. Model the evolution of the sick and immune populations over the next 120 days.

   Let $u(t)$, $s(t)$ and $i(t)$ represent the uninfected, sick and immune populations on day t. Then

   $$u(t+1) = 0.99u(t), \quad u(0) = 10000 \implies u(t) = 10000 \cdot 0.99^t$$

   The sick population is the sum of the previous 10 days' decrease in the at risk population:

   $$s(t) = \begin{cases} u(0) - u(t) = 10000(1 - 0.99^t) & \text{if } t \le 10 \\ u(t-10) - u(t) = 10000(0.99^{-10} - 1)0.99^t = 1057 \cdot 0.99^t & \text{if } t > 10 \end{cases}$$

   The immune population is the difference between these and the total population

   $$i(t) = 10000 - u(t) - s(t) = \begin{cases} 0 & \text{if } t \le 10 \\ 10000 - 11057 \cdot 0.99^t & \text{if } t > 10 \end{cases}$$

   After 120 days, we have

   $$u(120) = 2994, \qquad s(120) = 316, \qquad i(120) = 6690$$

2. Suppose 6000 vaccines are available. Discuss how these should be deployed. What should the goal be? Discuss the following strategies; for simplicity, assume the vaccines are 100% effective and work instantly.

   (a) Use all vaccine doses immediately.

   (b) Wait 30 days until some people are immune, then use all vaccines on the uninfected population.

   (c) Vaccinate 100 uninfected people per day.

   (d) Wait until there are 6000 uninfected people remaining, then vaccinate them all at once.

It is much easier to analyze this problem using a spreadsheet, though we can also do things analytically. Here are graphs of what happens under a non-intervention scenario and the first three vaccination campaigns.

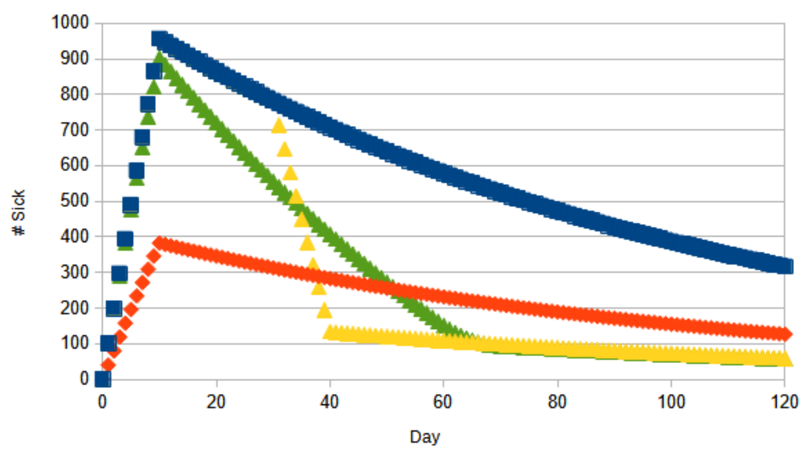

Exercises 3.3. 1. A cup of coffee is left outside on a warm day when the surrounding temperature is 90°F. Suppose the initial temperature of the coffee is 200°F and that its temperature after 2 minutes is 170°F. Find the temperature as a function of time.

2. Consider Corollary 3.14.

   (a) Check that $y(x) = M + (y_0 - M)e^{-kt}$ satisfies the differential equation.

   (b) A student believes that Theorem 3.6 is true. How would you convince them, *in the simplest possible way*, that the *only* solution to $y' = k(M - y)$ is as given in part (a)?

3. For both models in Example 3.16, what is the maximum growth rate of the yeast population, and at what time does it occur?

4. In Example 3.17, suppose the initial population is 3000 fish and 1000 fish are harvested per month, how long does it take for the population to recover to 6000 fish?

5. Plutonium-238 has a *half-life* of 88 years, meaning that after 88 years half of the isotope has decayed to another element (in this case uranium-234).

   (a) If you start with 100 grams of plutonium, find a model for how much remains after t years.

   (b) How long will take for the mass to decay to 10 grams, and at what rate will it be decaying?

6. (a) If public health officials wanted to eradicate the virus on page 42 completely by using all vaccine doses on one day, on which day should they act?

   (b) Find the number of uninfected people as a function of time under the 100 per day for 60 days vaccination scenario.

4 Sequences as Functions

We've seen many different types of function in this course and used them to model various situations. In practice, one is often faced with the opposite problem: given experimental data, what type of function should you try?

4.1 Polynomial Sequences: First, Second, and Higher Differences

To begin to answer this, first ask yourself, "What is a sequence?" Hopefully you have a decent intuitive idea already. More formally, a sequence a function whose domain is a set like the natural numbers, for example

$$f : \mathbb{N} \to \mathbb{R} : n \mapsto 3n^2 - 2$$

defines the sequence

$$\big(f(1), f(2), f(3), \ldots\big) = (1, 10, 25, 46, 73, \ldots)$$

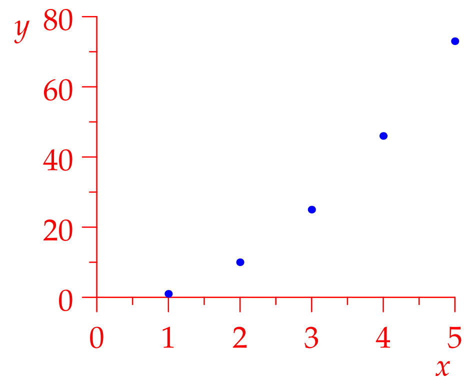

This is indeed the intuitive idea of a function to many grade-school students: continuity and domains including fractions or even irrational numbers are more advanced concepts.

Suppose instead that all you have is a data set

x	1	2	3	4	5
y	1	10	25	46	73

perhaps arising from an experiment. Could you recover the original function $y = f(x)$ directly from this data? You could try plotting data points as we've done, though it is hard to decide directly from the plot whether we should try a quadratic model, some other power function/polynomial, or perhaps an exponential. Of course, the physical source of real-world data might also provide clues.

A more mathematical approach involves considering how data values *change*:

	+1		+1		+1		+1		
x: 1		2		3		4		5	
y: 1		10		25		46		73	
	+9		+15		+21		+27		
		+6		+6		+6			

The first-differences in the x-values are constant whereas those for the y-values are increasing

$$(y_{n+1} - y_n) = (9, 15, 21, 27, \ldots)$$

You likely already notice the pattern: the sequence of first-differences is increasing *linearly* as the *arithmetic sequence*

$$y_{n+1} - y_n = 3 + 6n$$

To make this even clearer, note that the sequence of second-differences in the y-values is *constant* $(+6)$. These facts are huge clues that we expect a quadratic function.

But why? Well we can certainly check the following directly:

Linear Model If $f(n) = an + b$, then the sequence of first-differences is constant

$$f(n+1) - f(n) = a$$

Quadratic Model If $f(n) = an^2 + bn + c$, then the sequence of first-differences is linear and the second-differences are constant:

$$g(n) := f(n+1) - f(n) = 2an + a + b, \qquad g(n+1) - g(n) = 2a$$

The relationship between these results and the *derivative(s)* of the original function $f(x)$ should feel intuitive: what happens if you differentiate a quadratic twice?

Example 4.1. You are given the following data set

x	0	2	4	6	8	10
y	2	16	22	20	10	-8

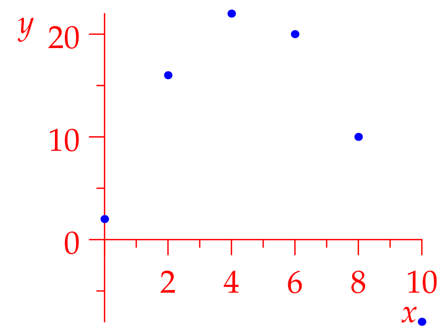

The x-values have constant first-differences while the y-values have constant second-differences

First-differences: $14, 6, -2, -10, -18$

Second-differences: $-8, -8, -8, -8$

We therefore suspect a quadratic model $y = f(n) = an^2 + bn + c$. Rather than using the above formulae, particularly since the x-differences are not 1, it is easier just to substitute:

$$2 = y(0) = c, \quad \begin{cases} 16 = f(2) = 4a + 2b + 2 \\ 22 = f(4) = 16a + 4b + 2 \end{cases} \implies \begin{cases} 2a + b = 7 \\ 8a + 2b = 10 \end{cases} \implies 4a = -4$$

whence $a = -1$, $b = 9$ and $c = 2$. A quadratic model is therefore

$$y = f(n) = -n^2 + 9n + c = -n^2 + 9n + 2$$

It is easily verified that the remaining data values satisfy this relationship.

There are at least two issues with our method:

1. The question we're answering is, "Find a quadratic model satisfying given data." Constant second-differences don't guarantee that only a quadratic model is suitable. For example,

   $$y = -n^2 + 9n + 2 + 297n(n-2)(n-4)(n-6)(n-8)(n-10)$$

   is a very complicated model satisfying the same data set!

2. It is very unlikely that experimental data will fit such precise patterns (why not?). However, if the differences are *close* to satisfying such patterns, then you should feel confident that a linear/quadratic model is a good choice.

Example 4.2. Given the data set

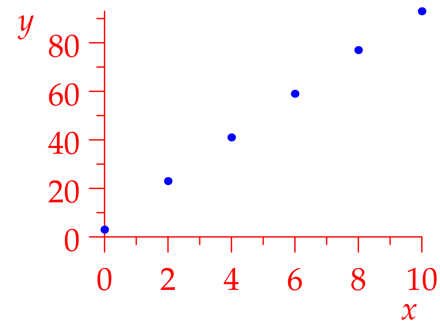

x	0	2	4	6	8	10
y	3	23	41	59	77	93

with sequences of first- and second-differences

First-differences: 20, 18, 18, 18, 16

Second-differences: $-2, 0, 0, -2$

do you think a linear or quadratic model would be superior?
If you wanted a linear model, you'd likely be inclined to try $f(x) = 9x + b$ for some constant b. Here are two options:

1. $f(x) = 9x + 5$ fits the middle four data values perfectly, but as a predictor is too large at the endpoints: $f(0) = 5 > 3$ and $f(10) = 95 > 93$.

2. $f(x) = 9x + 5 - \frac{2}{3}$ doesn't pass through any of the data values but seems to reduce the net error to zero:

x	0	2	4	6	8	10
$f(x) - y$	$-\frac{4}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	$\frac{2}{3}$	$-\frac{4}{3}$

$$\implies \sum_x f(x) - y = 0$$

Neither model is perfect, but then this is what you expect with real-world data!

Exercises 4.1. 1. For each data set, find a function $y = f(x)$ modelling the data.

(a)

x	2	4	6	8
y	-1	2	7	14

(b)

x	2	5	8	11	14
y	-6	-15	-6	21	66

(c)

x	0	6	9	15
y	3	15	21	33

(*Be careful with (c): the x-differences aren't constant!*)

2. Suppose a table of data values containing (x_0, y_0) has constant first-differences in both variables

$$\Delta x = x_{n+1} - x_n = a, \quad \Delta y = b$$

Find the equation of the linear function $y = f(x)$ through the data.

3. What relationship do you expect to find with the sequential differences of a cubic function $f(n) = an^3 + bn^2 + cn + d$? What about a degree-$m$ polynomial $f(n) = an^m + bn^{m-1} + \cdots$?

4. If $f(n) = an^2 + bn + c$ is a quadratic model for the data in Example 4.2 with constant second-differences -1, show that $a = -\frac{1}{8}$. What might be reasonable values for b, c?

5. (Hard) Suppose $f(x)$ is a twice-differentiable function and $h > 0$ is constant. Use the mean value theorem from calculus to explain the following.
   (a) First-differences $f(x+h) - f(x)$ are proportional to $f'(\xi)$ for some $\xi \in (x, x+h)$.
   (b) Second-differences satisfy $\big(f(x+2h) - f(x+h)\big) - \big(f(x+h) - f(x)\big) = f''(\xi)h\alpha$ for some ξ between x and $x+h$ and some α. Why is it unlikely that α is constant?

4.2 Exponential, Logarithmic & Power Sequences

To observe relationships between data values, you might also have to consider *ratios* between successive terms or skip values.

Example 4.3. From a first glance at the given data, it is hard to decide whether an exponential or a quadratic (or higher degree polynomial) model is more suitable. If we try to apply the constant-difference method, we don't seem to get anything helpful:

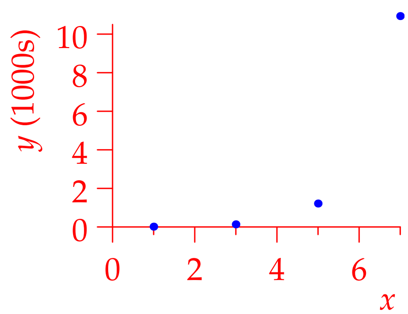

x	1	$\xrightarrow{+2}$	3	$\xrightarrow{+2}$	5	$\xrightarrow{+2}$	7
y	15	$\xrightarrow{+120}$	135	$\xrightarrow{+1080}$	1215	$\xrightarrow{+9720}$	10935
			$\xrightarrow{+960}$		$\xrightarrow{+8640}$		

By the time we're looking at second-differences, any conclusion would be very weak since we only have two data values!

If instead we think about *ratios* of y-values, then a different pattern emerges:

x	1	$\xrightarrow{+2}$	3	$\xrightarrow{+2}$	5	$\xrightarrow{+2}$	7
y	15	$\xrightarrow{\times 9}$	135	$\xrightarrow{\times 9}$	1215	$\xrightarrow{\times 9}$	10935

The question remains: what type of function scales its output by 9 when 2 is added to its input: $f(x+2) = 9f(x)$? This is a function that converts *addition* to *multiplication*: an exponential! If we try $y = f(x) = ba^x$ for some constants a, b, then

$$f(x+2) = ba^{x+2} = ba^2b^x = a^2 f(x)$$

from which a suitable model is $y = 5 \cdot 3^x$.

We can see the pattern in the example more generally:

Exponential Model If $f(x) = ba^x$, then adding a constant to x results in

$$f(x+k) = ba^{x+k} = a^k f(x)$$

If x-values have constant differences $(+k)$, then y-values will be related by a constant *ratio* $(\times a^k)$. You might remember this as 'addition–product' or 'arithmetic–geometric.'

Such a simple pattern is often disguised:

- Complete data might not be given so you might have to skip some data values to see a pattern. For example, if our original data was

x	1	3	4	5	7
y	15	135	405	1215	10935

  then the x-values are not in a strictly arithmetic sequence.

- As in Example 4.2, real-world/experimental data will only *approximately* exhibit such patterns.

Example 4.4. A population of rabbits is measured every two months resulting in the data set

t	0	2	4	6	8	10
P	5	7	10	14	19	28

The data seems very close to being quadratic; consider the first and second sequences of P-differences

$$\Delta P = (2, 3, 4, 5, 9), \qquad \Delta\Delta P = (1, 1, 1, 4)$$

However, the last difference doesn't fit the pattern. Instead, the fact that we expect an exponential model is buried in the experiment: the data is measuring population growth! We therefore instead consider the ratios of P-values:

		$+2$		$+2$		$+2$		$+2$		$+2$	
t	0		2		4		6		8		10
P	5		7		10		14		19		28
		$\times 1.4$		$\times 1.43$		$\times 1.4$		$\times 1.36$		$\times 1.47$	

The ratios are very close to being constant, whence an exponential model is suggested! To exactly match the first and last data values, we could take the model

$$P(t) \approx 5\left(\frac{28}{5}\right)^{\frac{t}{10}}$$

t	0	2	4	6	8	10
P	5	7.057	9.960	14.057	19.839	28

Only $P(8)$ doesn't match when we take rounding to the nearest integer into account.

We've seen that addition-addition corresponds to a linear model and that addition-multiplication to an exponential. There are two other natural combinations.

Logarithms These operate exactly as exponentials but in reverse. If $f(n) = \log_a x + b$, then *multiplying* x by a constant results in a constant *addition/subtraction* to y:

$$f(kx) = \log_a(kx) + b = \log_a k + \log_a x + b = \log_a k + f(x)$$

This could be summarized as 'product–addition.'

Power Functions If $f(x) = ax^m$, then multiplying x by a constant will do the same to y

$$f(kx) = a(kx)^m = ak^m x^m = k^m f(x)$$

We have a 'product–product' relationship between successive terms.

Examples 4.5. Find the patterns in the following data and suggest a model $y = f(x)$ in each case.

x	6	18	54	162
y	1	2	3	4

x	3	6	9	12
y	135	1080	3645	8640

The sequential approach in this chapter is a form of *discrete calculus*: using a pattern of *differences* to predict the original function is similar to how we use knowledge of a derivative $f'(x)$ to find $f(x)$.

Example 4.6. Suppose $g(2) = 3$ and $g(4) = 9$. What do you think should be the value of $g(8)$?

It depends on the type of model you try.

1. For a linear (addition-addition) model we know that $\Delta x = 2$ corresponds to $\Delta y = 6$, so

$$g(8) = g(4 + 2\Delta x) = g(4) + 2\Delta y = 9 + 12 = 21$$

2. For an exponential (addition-product) model, $\Delta x = 2$ corresponds to a *y-ratio* $r_y = \frac{9}{3} = 3$, so

$$g(8) = r_y g(6) = r_y^2 g(4) = 9 \cdot 9 = 81$$

3. For a power (product-product) model, $r_x = 2$ corresponds to a $r_y = 3$, so

$$g(8) = g(2 \cdot 4) = g(4r_x) = r_y g(4) = 3 \cdot 9 = 27$$

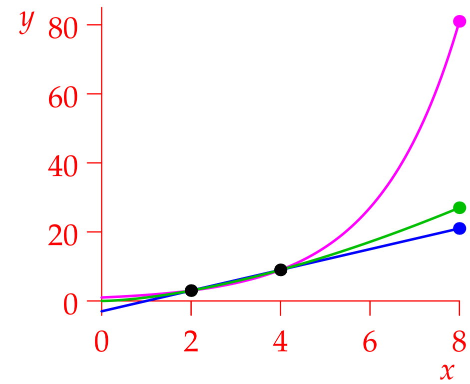

1. $g(x) = 3x - 3$
2. $g(x) = 3^{x/2}$
3. $g(x) = x^{\log_2 3}$

We do not need to calculate the models explicitly(!), though they are stated below the graph for convenience.

Exercises 4.2. 1. Find the patterns in the following data sets and use them to find a model $y = f(x)$.

(a)

x	0	1	2	3	4
y	80	120	180	270	405

(b)

x	2	4	8	10
y	1	16	256	625

(c)

x	1	3	5	7	9
y	15	5	19	57	119

(d)

x	1	3	4	6
y	1	36	216	7776

(e)

x	20	60	180	540
y	2	4	6	8

(f)

x	2	6	54	486	4374
y	2	4	8	12	16

2. Take logarithms of the power relationship $y = ax^m$. What is the relationship between $\ln y$ and $\ln x$? Use this to give another reason why the inputs and outputs of power functions satisfy a 'product–product' relationship.

3. How does our analysis of exponential functions change if we add a constant to the model? That is, how might you recognize a sequence arising from a function $f(x) = ba^x + c$?

4. Suppose $f(5) = 12$ and $f(10) = 18$. Find the value of $f(20)$ supposing $f(x)$ is a:

   (a) Linear function;

   (b) Exponential function;

   (c) Power function.

   If $f(20) = 39$, which of the three *models* do you think would be more appropriate?

4.3 Newton's Method

To finish our discussion of sequences we revisit a (hopefully) familiar technique for approximating solutions to equations. Variations of this approach have been in use for thousands of years.

Example 4.7. We motivate the method by considering an ancient method for approximating $\sqrt{2}$, known to the Babylonians 2500 years ago!
Suppose $x_n > \sqrt{2}$. Then $\frac{2}{x_n} < \frac{2}{\sqrt{2}} = \sqrt{2}$. It seems reasonable to guess that their average

$$x_{n+1} = \frac{1}{2}\left(x_n + \frac{2}{x_n}\right)$$

should be a more accurate approximation to $\sqrt{2}$. If start with an initial guess $x_0 = 2$, then we obtain the sequence

$$x_1 = \frac{1}{2}\left(2 + \frac{1}{2}\right) = \frac{3}{2}, \quad x_2 = \frac{17}{12} = 1.4166\ldots, \quad x_3 = \frac{577}{408} = 1.4142\ldots, \quad \ldots$$

This sequence certainly *appears* to be converging to $\sqrt{2}$...

Since it makes use of the average, this approach is sometimes called the *method of the mean*. It may be applied to any square-root $\sqrt{a}$ where $a > 0$: let $x_0 > 0$ and define,

$$x_{n+1} := \frac{1}{2}\left(x_n + \frac{a}{x_n}\right) \tag{$*$}$$

A rigorous proof that the sequence converges requires more detail than is appropriate for us (though see Exercise 3), but two observations should make it seem more believable:

1. *If* the sequence $(*)$ has a limit L, then the limit must satisfy

$$L = \frac{1}{2}\left(L + \frac{a}{L}\right) \implies 2L^2 = L^2 + a \implies L^2 = a \implies L = \sqrt{a}$$

   where we take the positive root since all terms x_n are plainly positive.

2. The iterations have a convincing *geometric* interpretation. The sequence of iterates can be found by repeatedly taking the tangent line to the curve $y = f(x) = x^2 - a$ and intersecting it with the x-axis. To see why, observe that the tangent line at x_n has equation

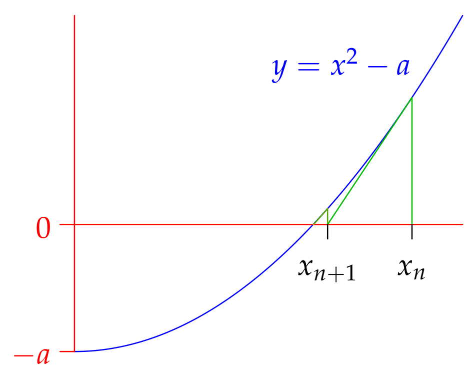

$$\begin{aligned} y &= f(x_n) + f'(x_n)(x - x_n) \\ &= x_n^2 - a + 2x_n(x - x_n) \\ &= 2x_n x - x_n^2 - a \end{aligned}$$

   which intersects the x-axis ($y = 0$) when

$$x = \frac{x_n^2 + a}{2x_n} = \frac{1}{2}\left(x_n + \frac{a}{x_n}\right) = x_{n+1}$$

This geometric idea generalizes...

Definition 4.8. Given a differentiable function $f(x)$ with non-zero derivative, the *Newton–Raphson iterates* of an initial value x_0 are defined by the recurrence formula

$$x_{n+1} := x_n - \frac{f(x_n)}{f'(x_n)}$$

Our two previous observations still hold:

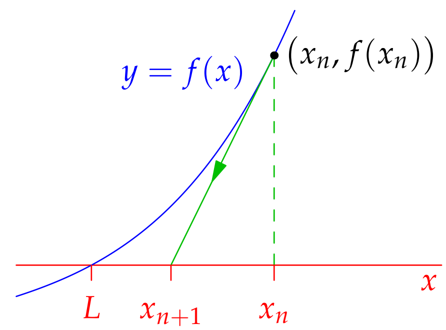

1. If $L = \lim_{n\to\infty} x_n$ exists and $f'(L) \neq 0$, then

$$L = L - \frac{f(L)}{f'(L)} \implies f(L) = 0$$

   That is, the limit L is a root of the function $f(x)$.

2. The tangent line at $(x_n, f(x_n))$ forms a right-triangle with base $x_n - x_{n+1}$ and height $f(x_n)$, from which its slope is

$$f'(x_n) = \frac{f(x_n)}{x_n - x_{n+1}}$$

   Rearranging this gives the formula $x_{n+1} = x_n - \frac{f(x_n)}{f'(x_n)}$.

Newton's method is particularly nice for polynomials with integer coefficients, since the iterates form a sequence of *rational numbers*. This approach was often used obtain rational approximations to irrational numbers before the advent of calculators.

Examples 4.9. 1. To find a root of $f(x) = x^4 + 4x - 6$, start with $x_0 = 2$ and iterate

$$x_{n+1} = x_n - \frac{x_n^4 + 4x_n - 6}{4x_n^3 + 4} = \frac{3(x_n^4 + 2)}{4(x_n^3 + 1)}$$

which yields the sequence (to 3 d.p.)

$$\left(2, \frac{3}{2}, \frac{339}{280}, \ldots\right) = (2, 1.5, 1.211, 1.121, 1.114, 1.114, \ldots)$$

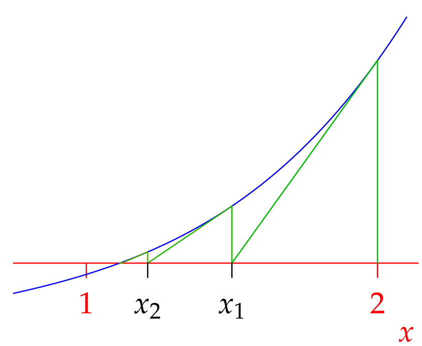

You can check with a calculator that 1.114 is approximately a root.

2. The irrational number $x = \sqrt{2} + \sqrt{3}$ is a root of the polynomial

$$f(x) = x^4 - 10x^2 + 1$$

By applying Newton's method with $x_0 = 3$, we obtain the sequence (to 3 d.p.)

$$x_{n+1} = x_n - \frac{x_n^4 - 10x_n^2 + 1}{4x_n^3 - 20x_n} = \frac{3x_n^4 - 10x_n^2 - 1}{4x_n(x_n^2 - 5)} \implies (x_n) = \left(3, \frac{19}{6} = 3.167, 3.147, \ldots\right)$$

Newton's method can be attempted for any differentiable function, though the sequence isn't guaranteed to converge: see for instance Exercise 5. You can find graphical interfaces online for this (for instance with Geogebra).

Exercises 4.3. 1. Use Newton's method to find a root of the given function to 4 decimal places.

(*Use a calculator, but explain what you are doing!*)

(a) $f(x) = x^3 - 4$ (b) $f(x) = 2x^3 + x - 1$ (c) $f(x) = e^x - \sqrt{x} - 2$

2. Use Newton's method to find a rational number approximation to $\sqrt[3]{2}$ in lowest terms $\frac{p}{q}$ where $10 < q < 100$.

3. Suppose you perform Newton's method for the function $f(x) = x^2 - 2$ starting with some positive $x_0 > 0$.

   (a) If $x_n > 0$, show that $x_{n+1} - \sqrt{2} = \frac{1}{2x_n}(x_n - \sqrt{2})^2 = \left(\frac{1}{2} - \frac{1}{2\sqrt{2}x_n}\right)(x_n - \sqrt{2})$.

   (b) Explain why $\left|x_n - \sqrt{2}\right| < \frac{1}{2^n}\left|x_0 - \sqrt{2}\right|$. Hence conclude that the sequence of iterates (x_n) converges to $\sqrt{2}$.

4. We might consider a *method of the mean* for approximating $\sqrt[3]{2}$: given x_0, define

$$x_{n+1} = \frac{1}{2}\left(x_n + \frac{2}{x_n^2}\right)$$

   (a) If the sequence (x_n) converges, show that its limit is $\sqrt[3]{2}$.

   (b) If $x_n > \sqrt[3]{2}$, show that $\frac{2}{x_n^2} < \sqrt[3]{2}$.

   (c) Let $x_0 = 1$. Compute x_1 and x_2. Compare these with the values obtained using Newton's method for the function $f(x) = x^3 - 2$ with the same initial condition $x_0 = 1$.

5. Let $f(x) = x^3 - 5x$.

   (a) What happens if you apply Newton's method to this function with initial condition $x_0 = 1$? Draw a picture to illustrate.

   (b) (Just for fun!) Investigate what happens for other values of x_0. Can you make any conjectures? Is is possible for x_0 to be *positive* and yet for $x_n \to -\sqrt{5}$? Can you make any sense of what happens if $1 < x_0 < \sqrt{\frac{5}{3}}$?

5 Regression Models

We've studied several types of function and seen how to spot whether a given data set might suit a particular model. To get further with this analysis, we need a method for comparing how bad a particular model is for given data.

5.1 Best-fitting Lines and Linear Regression

We start with an example of some data which appears reasonably linear.

Example 5.1. At t p.m., a trail-runner's GPS locator says that they've travelled y miles along a trail;

t_i	1	2	3	5
y_i	4	8	10	21

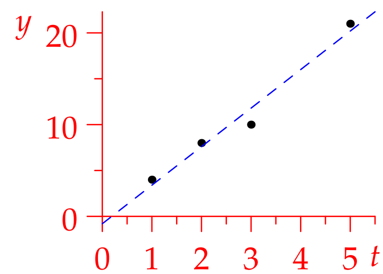

We'd like a simple model for how far the runner has travelled as a function of t. We might use this to predict where they would be at a given time; say at 6 p.m., or at 2 p.m. if they were to attempt the trail on another day.

By plotting the points, the relationship looks to be approximately[12] linear: $y \approx mt + c$. What is the *best* choice of line, and how should we find the coefficients m, c?

What might be good criteria for choosing our line? What should we mean by *best*? Plainly, we want the points to be close to the line, but measured how? What use do we want to make of the approximating line?

Here are three candidate lines plotted with the data set: of the choices, which seems best and why?

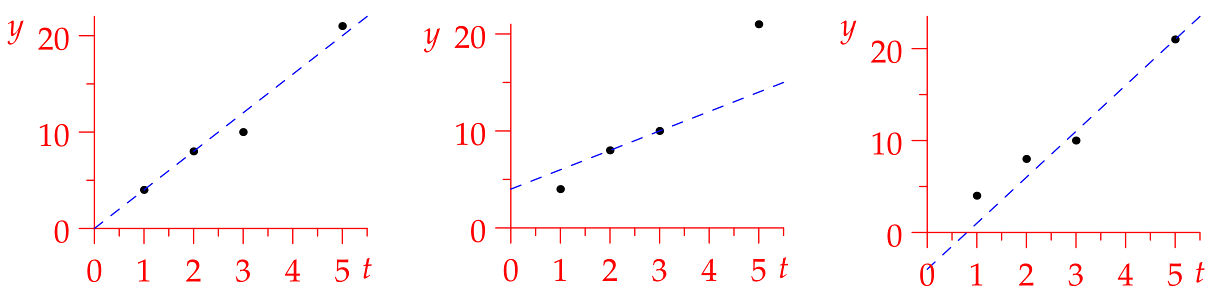

$y = 4t$ $\qquad$ $y = 2t + 4$ $\qquad$ $y = 5t - 4$

Since we want our model to *predict* the hiker's location $y \approx \hat{y} = mt + c$ at a given time t, we'd like our model to minimize *vertical* errors $\hat{y}_i - y_i$. We've computed these in the table; since a positive error is as bad as a negative, we make all the errors positive. It therefore seems reasonable to claim that the first line is the best choice of the three.

	t_i	1	2	3	5
	y_i	4	8	10	21
$y = 4t$	$\lvert\hat{y}_i - y_i\rvert$	0	0	2	1
$y = 2t + 4$	$\lvert\hat{y}_i - y_i\rvert$	2	0	0	7
$y = 5t - 4$	$\lvert\hat{y}_i - y_i\rvert$	3	2	1	0

But can we do better?

[12] Why should we not expect the distance traveled by the hiker to be perfectly linear?

We need a sensible definition of *best-fitting line* for a given data set. One possibility is to minimize the sum of the vertical errors:

$$\sum_{i=1}^{n} |\hat{y}_i - y_i|$$

For reasons of computational simplicity, uniqueness, statistical interpretation, and to discourage large individual errors, we *don't* do this! The standard approach is instead to minimize the *sum of the squared errors.*

Definition 5.2. Let (t_i, y_i) be data points with at least two distinct t-values. Let $\hat{y} = mt + c$ be a linear predictor (model) for y given t.

- The i^{th} *error* in the model is the difference $e_i := \hat{y}_i - y_i = mt_i + c - y_i$.
- The *regression line* or *best-fitting least-squares line* is the function $\hat{y} = mt + c$ which minimizes the sum $S := \sum e_i^2 = \sum(\hat{y}_i - y_i)^2$ of the *squares* of the errors.

Having at least two distinct t-values (some $t_i \neq t_j$) is necessary for the regression line to be unique.

Example (5.1, cont). Suppose the predictor was $\hat{y} = mt + c$. We expand the table

t_i	1	2	3	5
y_i	4	8	10	21
$\hat{y}_i$	$m+c$	$2m+c$	$3m+c$	$5m+c$
e_i	$m+c-4$	$2m+c-8$	$3m+c-10$	$5m+c-21$

Our goal is to minimize the function

$$S(m,c) = \sum e_i^2 = (m+c-4)^2 + (2m+c-8)^2 + (3m+c-10)^2 + (5m+c-21)^2$$

This is easy to deal with if we invoke some calculus. If (m,c) minimizes $S(m,c)$, then the first derivative tests says that the (partial) derivatives of S must be zero.

- Keep c constant and differentiate with respect to m:

$$\begin{aligned}\frac{\partial S}{\partial m} &= 2(m+c-4) + 4(2m+c-8) + 6(3m+c-10) + 10(5m+c-21) \\ &= 2\Big[39m + 11c - 155\Big]\end{aligned}$$

- Keep m constant and differentiate with respect to c:

$$\begin{aligned}\frac{\partial S}{\partial c} &= (m+c-4) + (2m+c-8) + (3m+c-10) + (5m+c-21) \\ &= 11m + 4c - 43\end{aligned}$$

The regression line is found by solving a pair of simultaneous equations

$$\begin{cases} 39m + 11c = 155 \\ 11m + 4c = 43 \end{cases} \implies m = \frac{21}{5},\ c = -\frac{4}{5} \implies \hat{y} = \frac{1}{5}(21t - 4)$$

By 6 p.m., we predict that the runner would have covered 24.4 miles. The sum of the squared errors for our regression line is $\sum e_i^2 = \sum |\hat{y}_i - y_i|^2 = 4.4$, compared to 5, 53 and 14 for our earlier options.

To obtain the general result for n data points, we return to our computations of the partial derivatives:

$$\frac{\partial S}{\partial m} = \sum \frac{\partial}{\partial m}(mt_i + c - y_i)^2 = 2\sum t_i(mt_i + c - y_i) = 2\Big[\big(\sum t_i^2\big)\, m + \big(\sum t_i\big)\, c - \sum t_i y_i\Big]$$
$$\frac{\partial S}{\partial c} = \sum \frac{\partial}{\partial c}(mt_i + c - y_i)^2 = 2\sum (mt_i + c - y_i) = 2\Big[\big(\sum t_i\big)\, m + nc - \sum y_i\Big]$$

These sums are often written using a short-hand notation for *average*:

$$\bar{t} = \frac{1}{n}\sum_{i=1}^{n} t_i, \qquad \overline{t^2} = \frac{1}{n}\sum_{i=1}^{n} t_i^2, \qquad \bar{y} = \frac{1}{n}\sum_{i=1}^{n} y_i^2, \qquad \overline{ty} = \frac{1}{n}\sum_{i=1}^{n} t_i y_i$$

Theorem 5.3 (Linear Regression). *Given n data points (t_i, y_i) with at least two distinct t-values, the best-fitting least-squares line has equation $\hat{y} = mt + c$, where m, c satisfy*

$$\begin{cases} \big(\sum t_i^2\big)\, m + \big(\sum t_i\big)\, c = \sum t_i y_i \\ \big(\sum t_i\big)\, m + nc = \sum y_i \end{cases} \quad \leftrightsquigarrow \quad \begin{cases} \overline{t^2} m + \bar{t} c = \overline{ty} \\ \bar{t} m + c = \bar{y} \end{cases}$$

This is a pair of simultaneous equations for the coefficients m, c, with solution

$$m = \frac{\overline{ty} - \bar{t}\bar{y}}{\overline{t^2} - \bar{t}^2}, \qquad c = \bar{y} - m\bar{t}$$

As the next section shows, having two distinct t-values guarantees a non-zero denominator $\overline{t^2} - \bar{t}^2$. The expression for c shows that the regression line passes through the data's *center of mass* $(\bar{t}, \bar{y})$.

Example 5.4. Five students' scores on two quizzes are given. If a student scores 9/10 on the first quiz, what might we expect them to score on the second?

Quiz 1	8	10	6	7	4
Quiz 2	10	7	5	8	6

To put the question in standard form, suppose Quiz 1 is the t-data and Quiz 2 the y-data. It is helpful to rewrite the data and add lines to the table so that we may more easily compute everything.

	Data					$\sum$	Average
t_i	8	10	6	7	4	35	7
y_i	10	7	5	8	6	36	7.2
t_i^2	64	100	36	49	16	265	53
$t_i y_i$	80	70	30	56	24	260	52

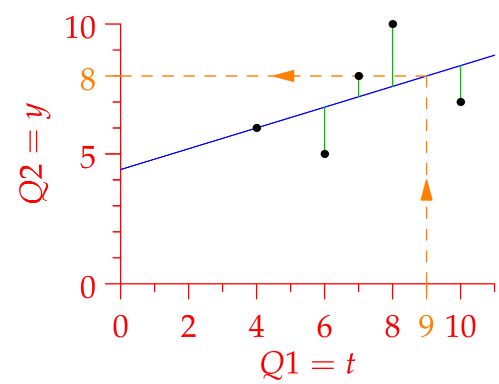

$$m = \frac{52 - 7 \times 7.2}{53 - 7^2} = \frac{1.6}{4} = 0.4, \qquad c = 7.2 - 0.4 \times 7 = 4.4$$

$$\implies \hat{y}(t) = \frac{2}{5}(t + 11)$$

This line which minimizes the sum of the squares of the vertical deviations. The prediction is that the hypothetical student scores $\hat{y}(9) = \frac{2}{5} \cdot 20 = 8$ on Quiz 2. Note that the predictor isn't symmetric: if we reverse the roles of t, y we don't get the same line!

Exercises 5.1. 1. Compute the sum of the absolute errors $\sum |\hat{y}_i - y_i|$ for the regression line and compare it to the sum of the absolute errors for $\hat{y} = 4t$: what do you notice?

2. Let $\hat{y} = mt + c$ be a linear predictor for the given data.

t_i	0	1	2	3
y_i	1	2	2	3

   (a) Compute the sum of squared-errors $S(m, c) = \sum e_i^2 = \sum |\hat{y}_i - y_i|^2$ as a function of m and c.

   (b) Compute the partial derivatives $\frac{\partial S}{\partial m}$ and $\frac{\partial S}{\partial c}$.

   (c) Find m and c by setting both partial derivatives to zero; hence find the equation of the regression line for these data.

   (d) Compare the sum of square errors S for the regression line with the errors if we use the simple predictor $y(t) = 1 + \frac{2}{3}t$ which passes through the first an last data points.

3. Consider Example 5.4.

   (a) Compute the sum of square-errors $S = \sum e_i^2 = \sum |\hat{y}_i - y_i|^2$ for the regression line.

   (b) Suppose a student was expected to score *exactly* the same on both quizzes; the predictor would be $\hat{y} = t$. What would the sum of squared-errors be in this case?

   (c) If a student scores 8/10 on *Quiz 2*, use linear regression to predict their score on Quiz 1. (*Warning: the answer is NOT* $\frac{5}{2} \cdot 8 - 11 = 9\ldots$)

4. Ten children had their heights (inches) measured on their first and second birthdays. The data was as follows.

1^{st} birthday	28	28	29	29	29	30	30	32	32	33
2^{nd} birthday	30	32	31	34	35	33	36	37	35	37

   Given this data, find a regression model and use it to predict the height at 2 years of a child who measures 32 inches at age 1.

   (*It is acceptable—and encouraged!—to use a spreadsheet to find the necessary ingredients. You can do this by hand if you like, but the numbers are large; it is easier with some formulæ from the next section.*)

5. (a) Let a, b be given. Find the value of y which minimizes the sum of squares

$$(y - a)^2 + (y - b)^2$$

   (b) For the data set $\{(t, y)\} = \{(1, 1), (2, 1), (2, 3)\}$, find the unique least-squares linear model for predicting y given t.
   (*Hint: think about part (a) if you don't want to compute*)

   (c) Show that there are *infinitely many* lines $\hat{y} = mt + c$ which minimize the sum of the absolute errors $\sum_{i=1}^{3} |\hat{y}_i - y_i|$.

5.2 The Coefficient of Determination

In the sense that it minimizes the sum of the squared errors $S = \sum e_i^2$, the linear regression model is as good as it can be—but *how* good? We could use S as a *quantitative* measure of the model's accuracy, but it doesn't do a good job at comparing the accuracy of models for *different* data sets. The standard approach to this problem relies the concept of variance.

Definition 5.5. The *variance* of data sequence $(y_1, \ldots, y_n)$ is the average of the squared deviations from their mean $\overline{y} = \frac{1}{n}\sum_{i=1}^{n} y_i$,

$$\operatorname{Var} y := \frac{1}{n}\sum_{i=1}^{n}(y_i - \overline{y})^2$$

The *standard deviation* is $\sigma_y := \sqrt{\operatorname{Var} y}$.

Variance and standard-deviation are measures of how data deviates from being constant.

Example 5.6. Suppose $(y_i) = (1, 2, 5, 4)$. Then

$$\overline{y} = \frac{1}{4}(1+2+5+4) = 3 \qquad \operatorname{Var} y = \frac{1}{4}\Big((-2)^2 + (-1)^2 + 2^2 + 1^2\Big) = \frac{5}{2} \qquad \sigma_y = \frac{\sqrt{10}}{2}$$

The square-root means that σ_y has the same units as y. Loosely speaking, a typical data value is expected to lie approximately $\sigma_y = \frac{1}{2}\sqrt{10} \approx 1.58$ from the mean $\overline{y} = 3$.

To obtain a measure for how well a regression line fits given data (t_i, y_i), we ask what *fraction* of the variance in y is explained by the model.

Definition 5.7. The *coefficient of determination* of a model $\hat{y} = mt + c$ is the ratio

$$R^2 := \frac{\operatorname{Var}\hat{y}}{\operatorname{Var} y}$$

Examples 5.8. We start by considering two extreme examples.

1. If the data were perfectly linear, then $y_i = mt_i + c$ for all i. The regression line is therefore $\hat{y} = mt + c$ and the coefficient of determination is precisely $R^2 = \frac{\operatorname{Var} y}{\operatorname{Var} y} = 1$. All the variance in the output y is explained by the model's transfer of the variance in the input t.

2. By contrast, consider the data in the table where we work out all necessary details to find the regression line:

$$m = \frac{\overline{ty} - \overline{t}\,\overline{y}}{\overline{t^2} - \overline{t}^2} = 0, \quad c = \overline{y} - m\overline{t} = 2$$

	data				average
t_i	0	0	2	2	$\overline{t} = 1$
y_i	1	3	1	3	$\overline{y} = 2$
t_i^2	0	0	4	4	$\overline{t^2} = 2$
t_iy_i	0	0	2	6	$\overline{ty} = 2$

The regression line is the *constant* $\hat{y} \equiv 2$, whence $\hat{y}$ has *no variance* and the coefficient of determination is $R^2 = 0$.
In this example, the regression model doesn't help explain the y-data in any way: the t-values have no obvious impact on the y-values.

In fact, the coefficient of determination always lies somewhere between these extremes $0 \le R^2 \le 1$: Exercise 6 demonstrates this and that the extreme situations are essentially those just encountered; in practice, therefore, $0 < R^2 < 1$. Before we revisit our examples from the previous section, observe that the average of the model's outputs $\hat{y}_i$ is the same as that of the original data:

$$\frac{1}{n}\sum_{i=1}^{n}\hat{y}_i = \frac{1}{n}\sum_{i=1}^{n}(mt_i + c) = m\bar{t} + c = \bar{y}$$

This makes computing the variance of $\hat{y}$ a breeze!

Example 5.1. Recall that $\hat{y} = \frac{1}{5}(21t - 4)$. Everything necessary is in the table

	data				average
t_i	1	2	3	5	$\bar{t} = 2.75$
y_i	4	8	10	21	$\bar{y} = 10.75$
$\hat{y}_i$	3.4	7.6	11.8	20.2	$\bar{\hat{y}} = 10.75$

$$\operatorname{Var} y = \frac{6.75^2 + 2.75^2 + 0.75^2 + 10.25^2}{4} = 39.6875$$

$$\operatorname{Var} \hat{y} = \frac{7.35^2 + 3.15^2 + 1.05^2 + 9.45^2}{4} = 38.5875$$

from which $R^2 = \frac{\operatorname{Var}\hat{y}}{\operatorname{Var}y} = \frac{3087}{3175} \approx 97.23\%$. The interpretation here is that the data is very close to being linear; the output y_i is very closely approximated by the regression model with approximately 97% of its variance explained by the model.

Example 5.4. This time $\hat{y} = \frac{2}{5}(t + 11)$.

	data					average
t_i	8	10	6	7	4	$\bar{t} = 7$
y_i	10	7	5	8	6	$\bar{y} = 7.2$
$\hat{y}_i$	7.6	8.4	6.8	7.2	6	$\bar{\hat{y}} = 7.2$

$$\operatorname{Var} y = \frac{2.8^2 + 0.2^2 + 2.2^2 + 0.8^2 + 1.2^2}{5} = 2.96$$

$$\operatorname{Var} \hat{y} = \frac{0.4^2 + 1.2^2 + 0.4^2 + 0^2 + 1.2^2}{5} = 0.64$$

from which $R^2 = \frac{\operatorname{Var}\hat{y}}{\operatorname{Var}y} = \frac{8}{37} \approx 21.62\%$. In this case the coefficient of determination is small, which indicates that the model does not explain much of the variation in the output.

The four examples are plotted below for easy visual comparison between the R^2-values.

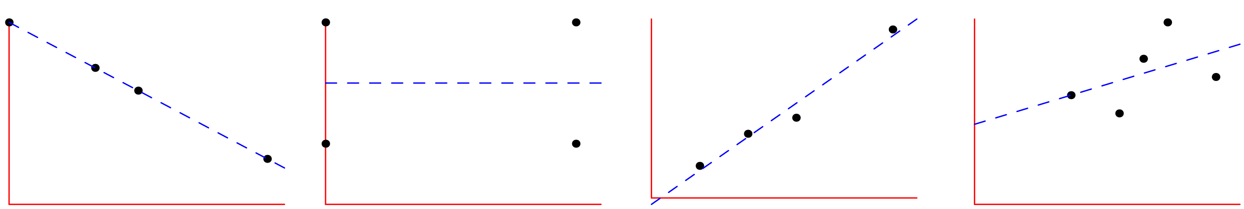

Perfect model $R^2 = 1$ Useless model $R^2 = 0$ Good model $R^2 = 0.97$ Poor model $R^2 = 0.22$

Efficient computation of R^2 If you want to compute by hand, our current process is lengthy and awkward. To obtain a more efficient alternative we first consider an alternative expression for the variance of any collection of data:

$$\operatorname{Var} x = \frac{1}{n}\sum(x_i - \bar{x})^2 = \frac{1}{n}\sum x_i^2 - \frac{2\bar{x}}{n}\sum x_i + \frac{\bar{x}}{n}\sum x_i = \overline{x^2} - \bar{x}^2$$

Plainly $\operatorname{Var} x \ge 0$ with equality if and only if all data values x_i are equal. The alternative expression $\overline{x^2} - \bar{x}^2$ justifies the uniqueness of the regression line in Definition 5.2 and Theorem 5.3.

Now expand the variance of the predicted outputs:

$$\operatorname{Var}\hat{y} = \frac{1}{n}\sum(\hat{y}_i - \overline{y})^2 = \frac{1}{n}\sum\big(mt_i + c - (m\bar{t} + c)\big)^2 = \frac{m^2}{n}\sum(t_i - \bar{t})^2 = m^2 \operatorname{Var} t$$

Putting these together, we obtain several equivalent expressions for the coefficient of determination:

$$R^2 = \frac{\operatorname{Var}\hat{y}}{\operatorname{Var} y} = m^2\frac{\operatorname{Var} t}{\operatorname{Var} y} = m^2\frac{\overline{t^2} - \bar{t}^2}{\overline{y^2} - \overline{y}^2} = \frac{(\overline{ty} - \bar{t}\overline{y})^2}{(\overline{t^2} - \bar{t}^2)(\overline{y^2} - \overline{y}^2)} \qquad (*)$$

Example 5.9. We do one more easy example with simple data $(t_i, y_i) : (1,4), (2,1), (3,2), (4,0)$.

data					average
t_i	1	2	3	4	$\bar{t} = \frac{10}{4}$
y_i	4	1	2	0	$\overline{y} = \frac{7}{4}$
t_i^2	1	4	9	16	$\overline{t^2} = \frac{15}{2}$
y_i^2	16	1	4	0	$\overline{y^2} = \frac{21}{4}$
t_iy_i	4	2	6	0	$\overline{ty} = 3$

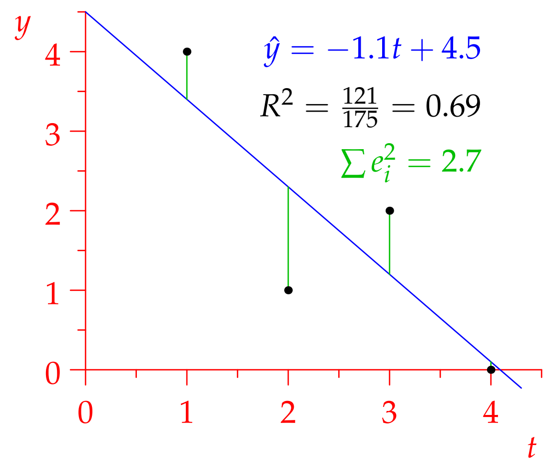

$$m = \frac{\overline{ty} - \bar{t}\overline{y}}{\overline{t^2} - \bar{t}^2} = \frac{3 - \frac{70}{4^2}}{\frac{15}{2} - \frac{100}{4^2}} = -\frac{11}{10} = -1.1$$

$$c = \overline{y} - m\bar{t} = \frac{7}{4} + \frac{11 \cdot 10}{10 \cdot 4} = \frac{9}{2} = 4.5$$

The regression line is $\hat{y} = -\frac{11}{10}t + \frac{9}{2} = -1.1t + 4.5$, and the coefficient of determination is

$$R^2 = m^2\frac{\overline{t^2} - \bar{t}^2}{\overline{y^2} - \overline{y}^2} = \frac{121}{100} \cdot \frac{\frac{15}{2} - \frac{100}{4^2}}{\frac{21}{4} - \frac{49}{4^2}} = \frac{121}{100} \cdot \frac{20}{35} = \frac{121}{175} = 69.1\%$$

The minimized square error is also easily computed:

$$\sum e_i^2 = \sum(\hat{y}_i - y_i)^2 = (3.4 - 4)^2 + (2.3 - 1)^2 + (1.2 - 2)^2 + (0.1 - 0)^2 = 2.7$$

Reversion to the Mean & Correlation By $(*)$, the regression model may be re-written in terms of the standard-deviation and R^2:

$$\hat{y}(t) = mt + c = \overline{y} + m(t - \bar{t}) = \overline{y} + \sqrt{R^2}\frac{\sigma_y}{\sigma_t}(t - \bar{t}) \implies \hat{y}(\bar{t} + \lambda\sigma_t) = \overline{y} + \lambda\sqrt{R^2}\sigma_y$$

Definition 5.10. The *correlation coefficient* is the value $r := \pm\sqrt{R^2}$ (sign equal to that of m).

An input λ standard-deviations above the mean ($t = \bar{t} + \lambda\sigma_t$) results in a prediction λr standard-deviations above the mean ($\hat{y} = \overline{y} + \lambda r\sigma_y$). Unless the data is perfectly linear, we have $R^2 < 1$; relative to the 'neutral' measure given by the standard-deviation a prediction $\hat{y}(t)$ is closer to the mean than the input t

$$\frac{|\hat{y}(t) - \overline{y}|}{\sigma_y} = r\frac{|\hat{t} - \bar{t}|}{\sigma_t} < \frac{|\hat{t} - \bar{t}|}{\sigma_t}$$

Example (5.9, cont). We compute the details. The correlation coefficient is $r = -\sqrt{R^2} \approx -0.832$; we say that the data is *negatively correlated*, since the output y seems to *decrease* as t increases. The standard deviations may be read off from the table:

$$\sigma_t = \sqrt{\operatorname{Var} t} = \sqrt{\overline{t^2} - \bar{t}^2} = \frac{\sqrt{5}}{2} \approx 1.118, \qquad \sigma_y = \sqrt{\operatorname{Var} y} = \sqrt{\overline{y^2} - \bar{y}^2} = \frac{\sqrt{35}}{4} \approx 1.479$$

The predictor may therefore be written (approximately)

$$\hat{y}(\bar{t} + \lambda\sigma_t) = \hat{y}(2.5 + 1.12\lambda) = \bar{y} + \lambda r \sigma_y = 1.75 - 1.23\lambda$$

As a sanity check,

$$\hat{y}(2.5 + 1.12) = \hat{y}(3.62) = -1.1 \times -3.98 + 4.5 = 0.52 = 1.75 - 1.23$$

Weaknesses of Linear Regression There are two obvious issues:

- Outliers massively influence the regression line. Dealing with this problem is complicated and there are a variety of approaches that can be used. It is important to remember that any approach to modelling, including our regression model, requires some *subjective choice.*
- If the data is not very linear then the regression model will produce a weak predictor. There are several ways around this as we'll see in the remaining sections: higher-degree polynomial regression can be performed, and data sometimes becomes more linear after some manipulation, say by an exponential or logarithmic function.

Exercises 5.2. 1. Suppose $(z_i) = (2, 4, 10, 8)$ is *double* the data set in Example 5.6. Find $\bar{z}, \operatorname{Var} z$ and σ_z. Why are you not surprised?

2. Use a spreadsheet to find R^2 for the predictor in Exercise 5.1.4. How confident do you feel in your prediction?

3. Find the standard deviations and correlation coefficients for the data in Examples 5.1 and 5.4.

4. The adult heights of men and women in a given population satisfy the following:

   Men: average 69.5 in, $\sigma = 3.2$ in. Women: average 63.7 in, $\sigma = 2.5$ in.

   The height of a father and his adult daughter have correlation coefficient 0.35. If a father's height is 72 in (mother's height unknown), how tall do you expect their daughter to be?

5. Suppose R^2 is the coefficient of determination for a linear regression model $\hat{y} = mt + c$. Use one of the alternative expressions for R^2 (page 59) to find the coefficient of determination for the reversed predictor $\hat{t}(y)$? Are you surprised?

6. Suppose that a data set $\{(t_i, y_i)\}_{1 \le i \le n}$ has at least two distinct t- and y-values (some $t_i \ne t_j$, etc.), that it has regression line $\hat{y} = mt + c$ and coefficient of determination R^2.

   (a) Show that $R^2 = 0 \iff m = 0$.
   (b) (Hard) Prove that the sum of squared errors equals $S = \sum_{i=1}^{n} e_i^2 = n(\operatorname{Var} y - \operatorname{Var} \hat{y})$.
   (c) Obtain the alternative expression $R^2 = 1 - \frac{S}{n \operatorname{Var} y}$. Hence conclude that $R^2 \le 1$, with equality if and only if the original data set is perfectly linear.

5.3 Matrix Multiplication & Polynomial Regression

In this section we consider how to find a best-fitting least-squares polynomial for given data. To see how to do this, it helps to rephrase the linear approach using matrices.[13]

We start by observing that the system of equations in Theorem 5.3 can be written in as a 2×2 matrix problem. For a data set with n pairs, the coefficients m, c satisfy

$$\begin{pmatrix} \sum t_i^2 & \sum t_i \\ \sum t_i & n \end{pmatrix} \begin{pmatrix} m \\ c \end{pmatrix} = \begin{pmatrix} \sum t_i y_i \\ \sum y_i \end{pmatrix}$$

This is nice because we can decompose the square matrix on the left as the product of a simple $2 \times n$ matrix and its transpose (switch the rows and columns);

$$\begin{pmatrix} \sum t_i^2 & \sum t_i \\ \sum t_i & n \end{pmatrix} = \begin{pmatrix} t_1 & t_2 & \cdots & t_n \\ 1 & 1 & \cdots & 1 \end{pmatrix} \begin{pmatrix} t_1 & 1 \\ t_2 & 1 \\ \vdots & \vdots \\ t_n & 1 \end{pmatrix} =: P^T P$$

We can also view the right side as the product of P^T and the column vector of output values y_i:

$$\begin{pmatrix} \sum t_i y_i \\ \sum y_i \end{pmatrix} = \begin{pmatrix} t_1 & t_2 & \cdots & t_n \\ 1 & 1 & \cdots & 1 \end{pmatrix} \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix} =: P^T \mathbf{y}$$

A little theory tells us that if *at least two* of the t_i are distinct, then the 2×2 matrix $P^T P$ is invertible;[14] there is a *unique* regression line whose coefficients may be found by taking the matrix inverse

$$\begin{pmatrix} m \\ c \end{pmatrix} = (P^T P)^{-1} P^T \mathbf{y} \implies \hat{y} = mt + c = (t \ \ 1) \begin{pmatrix} m \\ c \end{pmatrix} = (t \ \ 1)(P^T P)^{-1} P^T \mathbf{y}$$

We can also easily compute the vector of predicted values $\hat{y}_i = \hat{y}(t_i)$:

$$\hat{\mathbf{y}} = \begin{pmatrix} t_1 & t_2 & \cdots & t_n \\ 1 & 1 & \cdots & 1 \end{pmatrix} \begin{pmatrix} m \\ c \end{pmatrix} = P(P^T P)^{-1} P^T \mathbf{y}$$

and the squared error $\sum e_i^2 = \sum |\hat{y}_i - y_i|^2 = ||\hat{\mathbf{y}} - \mathbf{y}||^2$, which leads to an alternative expression for the coefficient of determination

$$R^2 = \frac{||\hat{\mathbf{y}}||^2 - n\overline{y}^2}{||\mathbf{y}||^2 - n\overline{y}^2}$$

where $||\mathbf{y}||$ is the *length* of a vector.

---

[13] Matrix computations are non-examinable. The purpose of this section is to be see how the regression may easily be automated and generalized by computer and to understand a little of how a spreadsheet calculates best-fitting curves of different types.

[14] For those who've studied linear algebra, P and $P^T P$ have the same null space and thus rank, since

$$P\mathbf{x} = \mathbf{0} \implies P^T P\mathbf{x} = \mathbf{0} \quad \text{and} \quad P^T P\mathbf{x} = \mathbf{0} \implies \mathbf{x}^T P^T P\mathbf{x} = \mathbf{0} \implies ||P\mathbf{x}|| = \mathbf{0} \implies P\mathbf{x} = \mathbf{0}$$

For linear regression, having at least two distinct t_i values means rank $P = 2$, whence $P^T P$ is invertible.

Examples 5.11. 1. We revisit the Example 5.9 in this language.

$$P = \begin{pmatrix} t_1 & 1 \\ t_2 & 1 \\ \vdots & \vdots \\ t_n & 1 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 2 & 1 \\ 3 & 1 \\ 4 & 1 \end{pmatrix} \implies P^T P = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 2 & 1 \\ 3 & 1 \\ 4 & 1 \end{pmatrix} = \begin{pmatrix} 30 & 10 \\ 10 & 4 \end{pmatrix}$$

from which

$$\begin{pmatrix} m \\ c \end{pmatrix} = (P^T P)^{-1} P^T \mathbf{y} = \begin{pmatrix} 30 & 10 \\ 10 & 4 \end{pmatrix}^{-1} \begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} 4 \\ 1 \\ 2 \\ 0 \end{pmatrix}$$

$$= \frac{1}{30 \cdot 4 - 10^2} \begin{pmatrix} 4 & -10 \\ -10 & 30 \end{pmatrix} \begin{pmatrix} 12 \\ 7 \end{pmatrix} = \frac{1}{20} \begin{pmatrix} 48 - 70 \\ -120 + 210 \end{pmatrix} = \frac{1}{10} \begin{pmatrix} -11 \\ 45 \end{pmatrix}$$

The prediction vector given inputs t_i is therefore

$$\hat{\mathbf{y}} = P \begin{pmatrix} m \\ c \end{pmatrix} = \frac{1}{10} \begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} -11 \\ 45 \end{pmatrix} = \frac{1}{10} \begin{pmatrix} 34 \\ 23 \\ 12 \\ 1 \end{pmatrix}$$

from which the coefficient of determination is, as before

$$R^2 = \frac{||\hat{\mathbf{y}}||^2 - 4\overline{y}^2}{||\mathbf{y}||^2 - 4\overline{y}^2} = \frac{\frac{1}{100}(34^2 + 23^2 + 12^2 + 1^2) - 4 \cdot \frac{7^2}{4^2}}{(4^2 + 1^1 + 2^2 + 0^2) - 4 \cdot \frac{7^2}{4^2}} = \frac{121}{175}$$

2. Given the data set $\{(3,1), (3,5), (3,6)\}$, we have $P = \begin{pmatrix} 3 & 1 \\ 3 & 1 \\ 3 & 1 \end{pmatrix}$ and $P^T P = \begin{pmatrix} 27 & 9 \\ 9 & 3 \end{pmatrix}$ which isn't invertible: $27 \cdot 3 - 9 \cdot 9 = 0$. The linear regression method doesn't work!

   It is easy to understand this from the picture. Since the three data points are vertically aligned, any line minimizing the sum of the squared errors must pass through the average $(3,4)$, though it could have *any* slope!

   This illustrates our fundamental assumption: linear regression requires at least two distinct t-values.

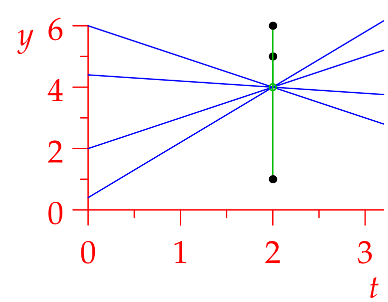

It is unnecessary ever to use the matrix approach for linear regression, though the method has significant advantages.

- Computers store and manipulate data in matrix format, so this method is computer-ready.
- Suppose you repeat an experiment several times, taking measurements y_i at times t_i. Since P depends only on the t-data, you need only compute the matrix $(P^T P)^{-1} P^T$ *once*, making computation of the regression line for repeat experiments very efficient.
- The method generalizes (easily for computers!) to polynomial regression…

Polynomial Regression

The pattern is almost identical when we use matrices; you just need to make the matrix P a little larger... We work through the approach for a quadratic approximation.

Suppose we have a data set $\{(t_i, y_i) : 1 \le i \le n\}$ and that we desire a quadratic polynomial predictor $\hat{y} = at^2 + bt + c$ which minimizes the sum of the squared vertical errors

$$S(a,b,c) = \sum_{i=1}^{n} e_i^2 = \sum_{i=1}^{n} (at_i^2 + bt_i + c - y_i)^2$$

This might look terrifying, but can be attacked exactly as before using differentiation: to minimize S, we need the derivatives of S with respect to the coefficients a, b, c to be zero.

$$\begin{cases} \dfrac{\partial S}{\partial a} = 2\sum at_i^4 + bt_i^3 + ct_i^2 - t_i^2 y_i = 0 \\ \dfrac{\partial S}{\partial b} = 2\sum at_i^3 + bt_i^2 + ct_i - t_i y_i = 0 \\ \dfrac{\partial S}{\partial c} = 2\sum at_i^2 + bt_i + c - y_i = 0 \end{cases} \iff \begin{cases} a\sum t_i^4 + b\sum t_i^3 + c\sum t_i^2 = \sum t_i^2 y_i \\ a\sum t_i^3 + b\sum t_i^2 + c\sum t_i = \sum t_i y_i \\ a\sum t_i^2 + b\sum t_i + cn = \sum y_i \end{cases}$$

As a system of equations for a, b, c this looks fairly nasty, but by rephrasing in terms of matrices, we see that it is exactly the same problem as before!

$$\begin{pmatrix} \sum t_i^4 & \sum t_i^3 & \sum t_i^2 \\ \sum t_i^3 & \sum t_i^2 & \sum t_i \\ \sum t_i^2 & \sum t_i & cn \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} \sum t_i^2 y_i \\ \sum t_i y_i \\ \sum y_i \end{pmatrix}$$

corresponds to

$$P^T P \begin{pmatrix} a \\ b \\ c \end{pmatrix} = P^T \mathbf{y} \quad \text{where} \quad P = \begin{pmatrix} t_1^2 & t_1 & 1 \\ \vdots & \vdots & \vdots \\ t_n^2 & t_n & 1 \end{pmatrix} \quad \text{and} \quad \mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_n \end{pmatrix}$$

The only change is that P is now an $n \times 3$ matrix so that $P^T P$ is 3×3. Analogous to the linear situation, provided at least three of the t_i are distinct, the matrix $P^T P$ is invertible and there is a unique least-squares quadratic minimizer

$$\hat{y} = at^2 + bt + c = \begin{pmatrix} t^2 & t & 1 \end{pmatrix} \begin{pmatrix} a \\ b \\ c \end{pmatrix} = \begin{pmatrix} t^2 & t & 1 \end{pmatrix} (P^T P)^{-1} P^T \mathbf{y}$$

The predictions $\hat{y}_i = \hat{y}(t_i)$ therefore form a vector $\hat{\mathbf{y}} = P \left(\begin{smallmatrix} a \\ b \\ c \end{smallmatrix}\right) = P(P^T P)^{-1} P^T \mathbf{y}$, and the coefficient of determination may be computed as before.

$$R^2 = \frac{||\hat{\mathbf{y}}||^2 - n\overline{y}^2}{||\mathbf{y}||^2 - n\overline{y}^2}$$

The method generalizes in the obvious way: if you want a cubic minimizer, give P an extra column of *cubed* t_i-terms! This would be hard work by hand, but is standard fodder for computers: this isn't a linear algebra class, so don't try to invert a 3×3 matrix!

Example 5.12. We are given data $\{(t_i, y_i)\} = \{(1,2),(2,5),(3,7),(4,4)\}$.

1. For the best-fitting linear model, we use the same P (and thus P^TP) from the previous example:

$$\begin{pmatrix} m \\ c \end{pmatrix} = (P^TP)^{-1}P^T\mathbf{y} = \begin{pmatrix} 30 & 10 \\ 10 & 4 \end{pmatrix}^{-1} \begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} 2 \\ 5 \\ 7 \\ 4 \end{pmatrix} = \frac{1}{10}\begin{pmatrix} 2 & -5 \\ -5 & 15 \end{pmatrix}\begin{pmatrix} 49 \\ 18 \end{pmatrix} = \begin{pmatrix} 0.8 \\ 2.5 \end{pmatrix}$$

which yields $\hat{y}(t) = 0.8t + 2.5$. The predicted values and coefficient of determination are then

$$\hat{\mathbf{y}} = \begin{pmatrix} 1 & 2 & 3 & 4 \\ 1 & 1 & 1 & 1 \end{pmatrix}^{\!T}\!\!\!\!\!\! \begin{pmatrix} 0.8 \\ 2.5 \end{pmatrix} = \begin{pmatrix} 3.3 \\ 4.1 \\ 4.9 \\ 5.7 \end{pmatrix} \qquad R^2 = \frac{84.2 - 81}{94 - 81} \approx 0.2462$$

The linear model predicts only 24.6% of the variance in the output; not very accurate.

2. For a quadratic model; all that changes is the matrix P

$$P = \begin{pmatrix} 1 & 1 & 1 \\ 4 & 2 & 1 \\ 9 & 3 & 1 \\ 16 & 4 & 1 \end{pmatrix} \implies P^TP = \begin{pmatrix} 1 & 4 & 9 & 16 \\ 1 & 2 & 3 & 4 \\ 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 & 1 \\ 4 & 2 & 1 \\ 9 & 3 & 1 \\ 16 & 4 & 1 \end{pmatrix} = \begin{pmatrix} 354 & 100 & 30 \\ 100 & 30 & 10 \\ 30 & 10 & 4 \end{pmatrix}$$

$$\implies \begin{pmatrix} a \\ b \\ c \end{pmatrix} = (P^TP)^{-1}P^T \begin{pmatrix} 2 \\ 5 \\ 7 \\ 4 \end{pmatrix} = \begin{pmatrix} 354 & 100 & 30 \\ 100 & 30 & 10 \\ 30 & 10 & 4 \end{pmatrix}^{-1} \begin{pmatrix} 149 \\ 49 \\ 18 \end{pmatrix} = \begin{pmatrix} -1.5 \\ 8.3 \\ -5 \end{pmatrix}$$

from which $\hat{y} = -1.5t^2 + 8.3t - 5$. To quantify its accuracy, compute the vector of predicted values $\hat{y}_i = \hat{y}(t_i)$ and the coefficient of determination:

$$\hat{\mathbf{y}} = P \begin{pmatrix} -1.5 \\ 8.3 \\ -5 \end{pmatrix} = \begin{pmatrix} 1.8 \\ 5.6 \\ 6.4 \\ 4.2 \end{pmatrix} \qquad R^2 = \frac{||\hat{\mathbf{y}}||^2 - 4\overline{y}^2}{||\mathbf{y}||^2 - 4\overline{y}^2} = \frac{93.2 - 81}{94 - 81} \approx 0.9385$$

The quadratic model is far superior to the linear, explaining 94% of the observed variance.

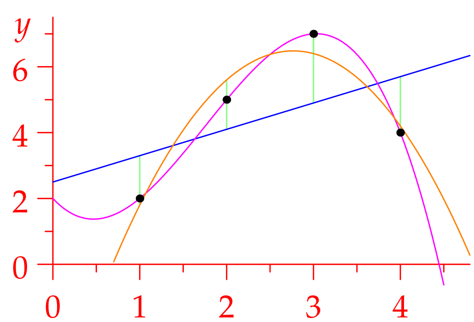

3. We can even find a cubic model (P is a 4×4 matrix!)

$$\hat{y} = \frac{1}{6}(-4t^3 + 21t^2 - 17t + 12)$$

The cubic passes through all four data points, there is *no error* and $R^2 = 1$.

For real-world data this is possibly *less useful* than the quadratic model—it certainly takes longer to find! More importantly, likely experimental error in the y-data has a strong effect on the 'perfect' model—we are, in effect, modelling *noise.* Do you expect $y(5)$ to be closer to -1 or -8?

Exercises 5.3. 1. Recall Example 4.2, with the following *almost* linear data set.

x	0	2	4	6	8	10
y	3	23	41	59	77	93

Find the best-fitting straight line for the data, then use a spreadsheet to find the best-fitting quadratic. Is the extra effort worth it?

2. You are given the following data consisting of measurements from an experiment recorded at times t_i seconds.

t_i	1	2	3	4	5	6	7	8	9	10
y_i	7	5	3	2	3	5	6	9	8	12

(a) Given the values

$$\sum t_i = 55, \quad \sum t_i^2 = 385, \quad \sum y_i = 60, \quad \sum t_i y_i = 385$$

find the best-fitting least-squares linear model for this data, and use it to predict $\hat{y}(13)$.

(b) Find the best-fitting quadratic model for the data: feel free to use a spreadsheet!

(c) The graphs below show the best-fitting least-squares linear, quadratic, cubic, quartic, and ninth-degree models and their coefficients of determination.

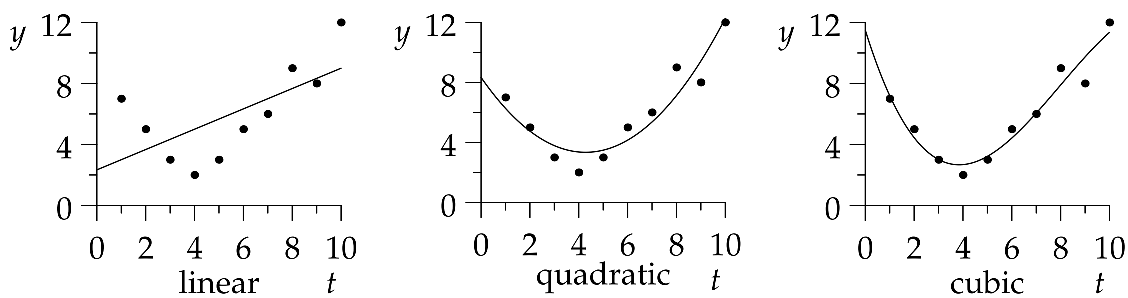

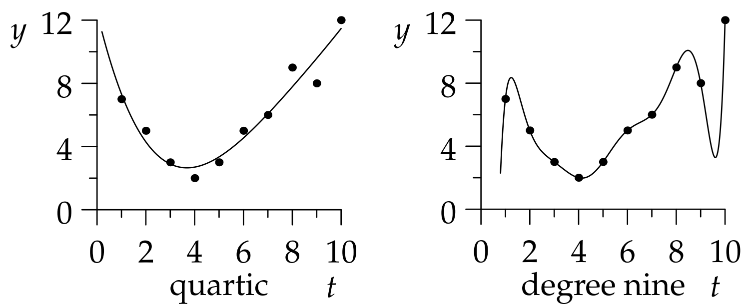

Degree	R^2
1	0.4264
2	0.8830
3	0.9319
4	0.9336
⋮	⋮
9	1

Which of these models would *you* choose for this data and why? What considerations would you take into account?

5.4 Exponential & Power Regression Models

If you suspect that your data would be better modelled by a non-polynomial function, there are several things you can try.

Minimizing the sum of squared-errors might be very difficult for non-polynomial functions because there is likely no simple tie-in with linear equations/algebra. Attempting this is likely to result in a horrible *non-linear* system for your coefficients which is difficult to analyze either theoretically or using a computer.[15]

Log Plots The most common approach when trying to fit an exponential model $\hat{y} = e^{mt+c}$ to data is to use a log plot: taking logarithms of both sides results in

$$\ln \hat{y} = mt + c$$

If we take $\hat{Y} := \ln \hat{y}$ as a new variable, the model is now a *straight-line*! The idea is then to use linear regression to find the coefficients $m, \ln a$.

Example (4.4, cont). Recall our earlier rabbit-population $P(t)$, repeated in the table below. We previously considered modelling this with an exponential function for two reasons:

1. We were told it was population data!
2. The *t-differences* are constant (2), while the *P-ratios* are approximately so (≈ 1.41).

t_i	0	2	4	6	8	10
P_i	5	7	10	14	19	28
$\ln P_i$	1.61	1.95	2.30	2.64	2.94	3.33

After constructing a log-plot, the relationship is much clearer:

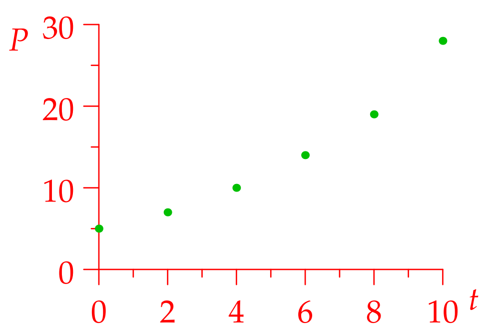

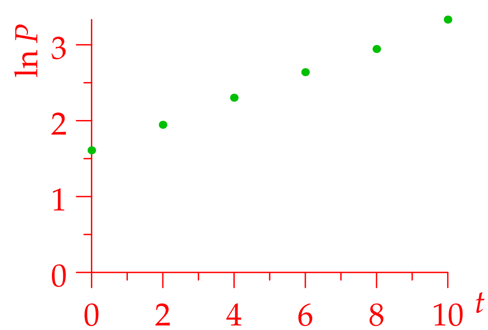

Since the relationship between t and $\ln P$ appears linear, we perform a linear regression calculation to find the best-fitting least-squares line for the $(t_i, \ln P_i)$ data.

[15] As an example of how horrific this is, suppose you want to minimize the sum of square-errors for data (t_i, y_i) using an exponential model $\hat{y}(t) = ae^{kt}$. The coefficients of our model, a, k should minimize

$$S(a,k) = \sum_{i=1}^{n} (ae^{kt_i} - y_i)^2$$

Differentiating this with respect to a, k and setting equal to zero results in

$$\begin{cases} \frac{\partial S}{\partial a} = 2\sum e^{kt_i}(ae^{kt_i} - y_i) = 0 \\ \frac{\partial S}{\partial k} = 2a\sum t_i e^{kt_i}(ae^{kt_i} - y_i) = 0 \end{cases} \implies \left(\sum y_i e^{kt_i}\right)\left(\sum t_i e^{2kt_i}\right) = \left(\sum e^{2kt_i}\right)\left(\sum t_i y_i e^{kt_i}\right)$$

where we substituted for a to obtain the last equation. Remember that this is an equation for k; if you think you can solve this easily, think again!

Everything necessary comes from extending the table.

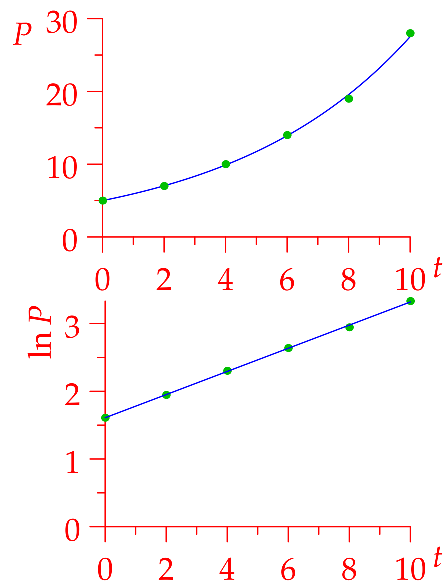

	Data						average
t_i	0	2	4	6	8	10	5
P_i	5	7	10	14	19	28	13.83
$\ln P_i$	1.61	1.95	2.30	2.64	2.94	3.33	2.46
t_i^2	0	4	16	36	64	100	36.67
$t_i \ln P_i$	0	3.89	9.21	15.83	23.56	33.32	14.30

$$m = \frac{\overline{t \ln P} - \bar{t} \cdot \overline{\ln P}}{\overline{t^2} - \bar{t}^2} = \frac{14.30 - 5 \cdot 2.46}{36.67 - 5^2} = 0.171$$

$$c = \overline{\ln P} - m\bar{t} = 3.46 - 0.171 \cdot 5 = 1.609$$

which yields the exponential model

$$\hat{P}(t) = e^{0.171t+1.609} = 4.998(1.186)^t$$

This is very close to the model $(5(1.188)^t)$ we obtained previously by pure guesswork. The approximate doubling time T for the population satisfies

$$e^{mT} = 2 \implies T = \frac{\ln 2}{m} = 4.06 \text{ months}$$

When using the log plot method, interpreting errors and the goodness of fit of a model is a little more difficult. Typically one computes the coefficient of determination R^2 of the *underlying linear model*: in our example,[16]

$$R^2 = m^2 \frac{\operatorname{Var} t}{\operatorname{Var} \ln P} = 99.3\%$$

It is important to appreciate that the log plot method does not treat all errors equally: taking logarithms tends to reduce error by a greater amount when the output y is large. This should be clear from the picture, and more formally by the mean value theorem: if $y_1 < y_2$, then there is some $\xi \in (y_1, y_2)$ for which

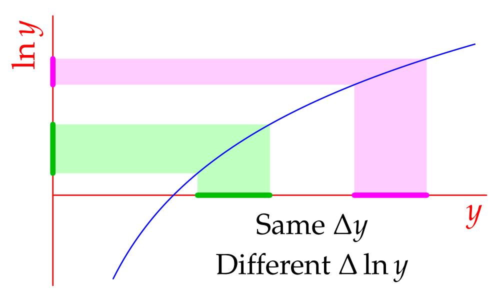

$$\ln y_2 - \ln y_1 = \frac{1}{\xi}(y_2 - y_1) < \frac{1}{y_1}(y_2 - y_1)$$

The log plot approach therefore places a higher emphasis on accurately matching data when the output y is *small*. This isn't such a bad thing since our intuitive view of error depends on the size of the data. For instance, misplacing a \$100 bill is annoying, but a \$100 mistake in escrow when buying a house is unlikely to concern you very much! Exponential data can more easily vary over large orders of magnitude than linear or quadratic data.

[16] This needs more decimal places of accuracy for the log-values than what's in our table!

Log-Log Plots If you suspect a *power function model* $\hat{y} = at^m$, then taking logarithms

$$\ln \hat{y} = m \ln t + \ln a$$

results in a linear relationship between $\ln y$ and $\ln t$. As before, we can apply a linear regression approach to find a mode; the goodness of fit is again described by the coefficient of determination of the underlying model.

Exercises 5.4. 1. You suspect a logarithmic model for a data set. Describe how you would approach finding a model in the context of this section.

2. The table shows the average weight and length of a fish species measured at different ages.

Age (years)	Length(cm)	Weight (g)
1	5.2	2
2	8.5	8
3	11.5	21
4	14.3	38
5	16.8	69
6	19.2	117
7	21.3	148
8	23.3	190
9	25.0	264
10	26.7	293
11	28.2	318
12	29.6	371
13	30.8	455
14	32.0	504
15	33.0	518
16	34.0	537
17	34.9	651
18	36.4	719
18	37.1	726
20	37.7	810

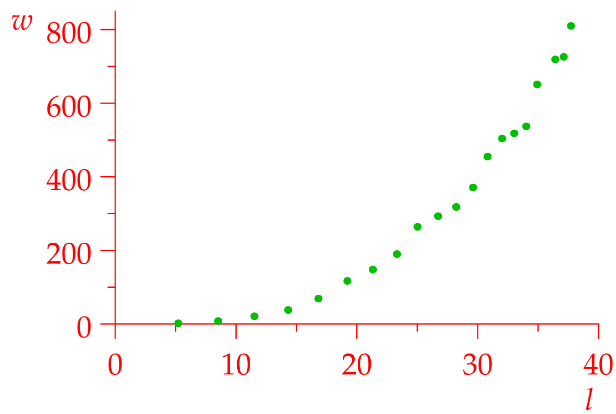

(a) Do you think an exponential model is a good fit for this data? Take logarithms of the weight values and use a spreadsheet to obtain a model $\hat{w}(\ell) = ae^{m\ell}$ where w, ℓ are the weight and length respectively.

(b) What happens if you try a log-log plot? Given what we're measuring, why do you expect a power model to be mode accurate?

3. Population data for Long Beach CA is given.

Using a spreadsheet or otherwise, find linear, quadratic, exponential and logarithmic regression models for this data.

Which of these models seems to fit the data best, and which would you trust to best predict the population in 2020?

Look up the population of Long Beach in 2020; does it confirm your suspicions? What do you think is going on?

Year	Years since 1900	Population
1900	0	2,252
1910	10	17,809
1920	20	55,593
1930	30	142,032
1940	40	164,271
1950	50	250,767
1960	60	334,168
1970	70	358,879
1980	80	361,498
1990	90	429,433
2000	100	461,522
2010	110	462,257

4. In the early 1600s, Johannes Kepler used observational data to derive his *laws of planetary motion*, the third of which relates the orbital period T of a planet (how long it takes to go round the sun) to its (approximate) distance r from the sun.

Planet	T (years)	r (millions km)
Mercury	0.24	58
Venus	0.61	110
Earth	1	150
Mars	1.88	230
Jupiter	11.9	780
Saturn	29.5	1400
Uranus	84	2900
Neptune	165	4500

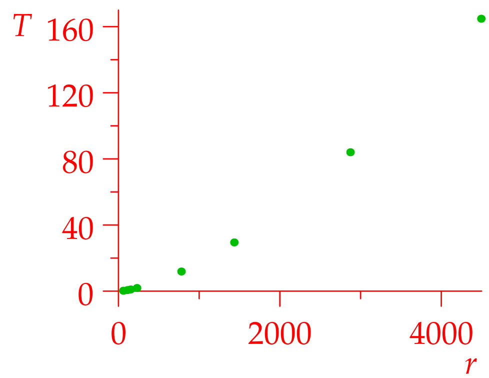

The table shows the data for all the planets. Use a spreadsheet to analyze this data and find a model relating T to r.

Kepler did not known about Uranus and Neptune and only had relative distances for the planets. Research the correct statement of Kepler's third law and compare it with your findings.