



CONTENTS

1	Data Tables and pandas	3
1.1	Data types	3
1.2	pandas	4
1.3	Reading a CSV with pandas	5
1.4	Looking at a Series	5
1.5	Rows and the index	7
1.6	Changing data	8
1.7	Derived columns	9
2	Atmospheric Pressure	13
2.1	Altitude and Atmospheric Pressure	15
2.2	How a Drinking Straw Works	16
2.2.1	The Longest Usable Straw	17
2.2.2	Millimeters Mercury	19
2.3	How a Siphon Works	20
2.4	How a Toilet Works	22
3	Complex Numbers	25
3.1	Definition	25
3.2	The Complex Plane	26
3.3	Adding and Subtracting	26
3.4	Multiplying	26
3.5	Conjugates and Division	27
3.6	Complex Numbers in Python	28

3.7	Why Complex Numbers Matter	28
4	Introduction to Sequences	31
4.1	Arithmetic sequences	32
4.1.1	Formulas for arithmetic sequences	32
4.2	Geometric sequences	35
4.2.1	Formulas for geometric sequences	36
5	Exponentials and Logarithms	39
5.1	Exponents	39
5.1.1	Identities for Exponents	39
5.2	Logarithms	41
5.2.1	Logarithms in Python	42
5.2.2	Logarithm Identities	43
5.2.3	Changing Bases	44
5.2.4	Natural Logarithm	44
5.2.5	Logarithms in Spreadsheets	44
5.3	Compound Interest	45
5.3.1	An example with annual interest payments	46
5.3.2	Exponential Growth	47
5.3.3	Sensitivity to interest rate	48
5.3.4	Continuous Compound Interest	49
5.3.5	Practice Exercises	49
5.4	Exponential Decay	50
5.4.1	Radioactive Decay	51
5.4.2	Model Exponential Decay	52
A	Answers to Exercises	55
	Index	61

Data Tables and pandas

Your school keeps a table of every student: name, grade, schedule. A streaming service keeps a table of every song you have played. Most of the data you will ever work with arrives as a table: a spreadsheet, a database, or a plain text file. In this chapter, you will learn pandas, the Python library that data scientists use to read, clean, and question tables.

Each column holds one attribute (like height or acreage), and each row holds one thing (like a person or a farm). Here is a table about houses:

property_id	bedrooms	square_meters	estimated_value
7927	3	121.4	\$ 294,393
9329	2	82.9	\$ 207,420

Usually one column is guaranteed to be different in every row. We call it the *primary key*, and it works like a student ID number. In this table, `property_id` is the primary key; every property has one, and no two properties have the same `property_id`.

1.1 Data types

Each column has a type, and the types match Python's types nicely.

	Type	Python type	Example
Here are some common datatypes:	Integer	int	910393
	Float	float	-23.19
	String	str	'Fred'
	Boolean	bool	False
	Date	datetime.date	2019-12-04
	Timestamps	datetime.datetime	2022-06-10T14:05:22Z

Sometimes it is OK to have values missing. For example, if you had a table of data about employees, maybe one of the columns would be retirement — a date that tells you when the person retired. People who had not yet retired would have no value in this column. We would say that they have *null* for retirement.

Sometimes there are constraints on what values can appear in the column. For example, if the column were height, it would make no sense to have a negative value.

Sometimes a column can only be one of a few values. For example, if you ran a bike rental shop, each bicycle's status would be "available", "rented", or "broken". Any other values

in that column would not be allowed. We often call these columns *categorical*.

1.2 pandas

Python programmers work with tables so often that they built a library just for them: pandas. You could loop through a table one row at a time, but for a table with millions of rows, that is slow. pandas lets you say things like “delete every row that has no value for height” in one line, and it runs fast.

In pandas, there are two datatypes that you use a lot:

- a Series is a single column of data.
- a DataFrame is a table of data: it has a Series for each column.

Worked example. Build a small DataFrame of test scores and find the average.

```
1 import pandas as pd
2
3 scores = pd.DataFrame({"name": ["Ana", "Ben", "Cy"], "score": [88, 92, 75]})
4 print(scores["score"].mean())
```

The output is 85.0. Square brackets pick out the score column (a Series), and mean() averages it.

Exercise 1 Highest Score

Using the scores DataFrame above, write code that prints the highest score.

Working Space

Answer on Page 55

In the digital resources, you will find bikes.csv, a table about the bikes in a rental shop. If you look at it in a text editor, it will start like this:

```
1 bike_id,brand,size,purchase_price,purchase_date,status
2 5636248,GT,57,277.99,1986-09-07,available
3 4156134,Giant,56,201.52,2005-01-09,rented
4 7971254,Cannondale,54,292.25,1978-02-28,available
```

```
5 | 3600023,Canyon,57,197.62,2007-02-15,broken
```

The first line is a header; it tells you the name of each column. Each line after that is one bike, with its values separated by commas. (Thus the name: CSV stands for “Comma-Separated Values”.)

1.3 Reading a CSV with pandas

Let us make a program that reads `bikes.csv` into a pandas dataframe. Create a file called `report.py` in the same folder as `bikes.csv`.

pandas treats one column as the primary key and calls it the *index*. When we read the file, we will tell pandas to use `bike_id` as the index.

If you ask a dataframe for its shape, it returns a tuple containing the number of rows and the number of columns. To confirm that we have actually read the data in, let us print those numbers. Add these lines to `report.py`:

```
1 import pandas as pd
2
3 # Read the CSV and create a dataframe
4 df = pd.read_csv('bikes.csv', index_col="bike_id")
5
6 # Show the shape of the dataframe
7 (row_count, col_count) = df.shape
8 print(f"*** Basics ***")
9 print(f"Bikes: {row_count:,}")
10 print(f"Columns: {col_count:,}")
```

Run it. You should see:

```
1 *** Basics ***
2 Bikes: 998
3 Columns: 5
```

The file has six columns, but the shape says five. The index does not count as a column.

1.4 Looking at a Series

Let us get the lowest, the highest, and the mean purchase price of the bikes. The purchase price is a series, and you can ask the dataframe for it. Add these lines to the end of your

program:

```
1 # Purchase price stats
2 print("\n*** Purchase Price ***")
3 series = df["purchase_price"]
4 print(f"Lowest:{series.min()}")
5 print(f"Highest:{series.max()}")
6 print(f"Mean:{series.mean():.2f}")
```

Now when you run it, you will see a few additional lines:

```
1 *** Purchase Price ***
2 Lowest:107.37
3 Highest:377.7
4 Mean:249.01
```

What are all the brands of the bikes? Add a few more lines to your program that shows how many of each brand:

```
1 # Brand stats
2 print("\n*** Brands ***")
3 series = df["brand"]
4 series_counts = series.value_counts()
5 print(f"{series_counts}")
```

Now when you run it, your report will include the number of bikes for each brand from most to least common:

```
1 *** Brands ***
2 brand
3 Canyon      192
4 BMC          173
5 Cannondale   170
6 Trek         166
7 GT           150
8 Giant        147
9 Name: count, dtype: int64
```

`value_counts` returns a Series. To format this better, we need to learn about accessing individual rows in a series.

1.5 Rows and the index

In a list, you ask for an item by its position. pandas does that with `iloc` (“integer location”). Add this to the end of your program:

```
1 # First bike
2 print("\n*** First Bike ***")
3 row = df.iloc[0]
4 print(f"{row}")
```

When you run it, you will see the attributes of the first row of data:

```
1 *** First Bike ***
2 brand                GT
3 size                57
4 purchase_price      277.99
5 purchase_date      1986-09-07
6 status              available
7 Name: 5636248, dtype: object
```

The row comes back as another Series.

The last line says that this row’s name (its index value) is 5636248. Usually we find rows by their index, not their position, using `loc`. For example, there is a bike with `bike_id` 2969341. Let us ask for its brand:

```
1 print("\n*** Some Bike ***")
2 brand = df.loc[2969341]['brand']
3 print(f"brand = {brand}")
```

Now, you will see the information about that bike:

```
1 *** Some Bike ***
2 brand = Cannondale
```

pandas has a few different ways of getting to that value. All of these get you the same thing:

```
1 brand = df.loc[2969341]['brand'] # Get row, then get value
2 brand = df['brand'][2969341]     # Get column, then get value
3 brand = df.loc[2969341, 'brand'] # One call with both row and value
```

1.6 Changing data

Real data is messy. Every bike should have a status, and it should be one of “available”, “rented”, or “broken”. Let us count each value that actually appears:

```

1 print("\n*** Status ***")
2 series = df["status"]
3 missing = series.isnull()
4 print(f"{missing.sum()} bikes have no status.")
5 series_counts = series.value_counts()
6 for value in series_counts.index:
7     print(f"{series_counts.loc[value]} bikes are \"{value}\"")

```

This will show you:

```

1 *** Status ***
2 7 bikes have no status.
3 389 bikes are "rented"
4 304 bikes are "broken"
5 296 bikes are "available"
6 1 bikes are "Flat tire"
7 1 bikes are "Available"

```

Somebody typed “Available” with a capital A, and somebody else typed “Flat tire.” Right after you read the CSV, fix the first problem:

```

1 mask = df['status'] == 'Available'
2 print(f"{mask}")
3 df.loc[mask, 'status'] = 'available'

```

When you run this, you will see that the mask is a series with bike_id as the index and False or True as the value, depending on whether the row’s status was equal to “Available”.

When you use loc with this sort of mask, you are saying “Give me all the rows for which the mask is True.” So, the assignment only happens in the one problematic row.

You do not need the mask variable. Here is the same fix in one line, plus one that turns Flat tire into broken:

```

1 df.loc[df['status'] == 'Available', 'status'] = 'available'
2 df.loc[df['status'] == 'Flat tire', 'status'] = 'broken'

```

Now those problems are gone:

```
1 7 bikes have no status.
2 389 bikes are "rented"
3 305 bikes are "broken"
4 297 bikes are "available"
```

What about the 7 bikes with no status? If we were sure they were available, we could fill in the blanks:

```
1 missing_mask = df['status'].isnull()
2 df.loc[missing_mask, 'status'] = 'available'
```

Or we could print their IDs so someone can go find them:

```
1 missing_mask = df['status'].isnull()
2 missing_ids = list(df[missing_mask].index)
3 print(f"These bikes have no status:{missing_ids}")
```

Instead, let us just keep the rows where the status is not null. The `~` flips every True to False and every False to True:

```
1 missing_mask = df['status'].isnull()
2 df = df[~missing_mask]
```

At the end of your program, write out the improved CSV:

```
1 df.to_csv('bikes2.csv')
```

Run the program and open `bikes2.csv` in a text editor.

1.7 Derived columns

Now add a new column, the age of each bike in days, so that `bikes2.csv` starts like this:

```
1 bike_id,brand,size,purchase_price,purchase_date,status,age_in_days
2 5636248,GT,57,277.99,1986-09-07,available,13061
3 4156134,Giant,56,201.52,2005-01-09,rented,6362
```

```
4 7971254,Cannondale,54,292.25,1978-02-28,available,16174
```

First problem: `purchase_date` looks like a date, but pandas read it as a string. Convert it by applying a function to every item in the Series. (Add `import datetime` at the top of your program.)

```
1 df['purchase_date'] = df['purchase_date'].apply(lambda s:
↳ datetime.date.fromisoformat(s))
```

(With pandas, there is often more than one way to do things. pandas has a `to_datetime` function that converts every entry in a sequence to a datetime object. Here is another way to convert the string column into a date column:

```
1 df['purchase_date'] = pd.to_datetime(df['purchase_date']).dt.date
```

You can look up `dt` and `date` if you are curious.)

Now, we can use the same trick to create a new column with the age in days:

```
1 today = datetime.date.today()
2 df['age_in_days'] = df['purchase_date'].apply(lambda d: (today - d).days)
```

Run it, and `bikes2.csv` will have an `age_in_days` column. (Your numbers will be bigger than ours; the bikes keep getting older.)

Exercise 2 Big bikes

Right after reading `bikes.csv`, how many bikes have a size of 57 or more? Hint: `df['size'] >= 57` is a mask, and adding up a mask counts its True values.

Working Space

Answer on Page 55

Exercise 3 The fanciest bike

Which bike has the highest purchase

Working Space

price? Print its `bike_id` and `brand`.
Hint: a Series has a method `idxmax` that returns the index of its largest value.

Answer on Page 55

CHAPTER 2

Atmospheric Pressure

Right now, the air is pressing on the palm of your hand with a force of about 1,000 newtons, the weight of a 100 kg sack of rocks. You do not feel a thing. This chapter explains why, and how that invisible push lifts juice up a straw, empties a fish tank, and flushes a toilet.

The air you breathe is a blend of gases:

1. 78% nitrogen in the form of N_2
2. 21% oxygen in the form O_2
3. 1% other gases (mostly argon)

Fill a balloon with helium (He), and the helium pushes outward on every bit of the balloon's skin. That push, spread over an area, is *pressure*: force divided by area. Force is measured in newtons, and pressure in *pascals*. A pascal is 1 newton per square meter.

The air outside is pushing back on the balloon, too. This is *atmospheric pressure* (also called *barometric pressure*). It comes from gravity pulling down on all the air above us.




Figure 2.1: Air pushes both ways on a balloon.

Picture a square meter of ground at sea level, and the column of air above it reaching all the way to the top of the atmosphere.




Figure 2.2: The column of air above you has a lot of weight.

The air in that column has a mass of about 10,340 kg, more than a loaded school bus. Each kilogram weighs 9.8 N, so the column pushes down on its square meter with $10,340 \times 9.8 \approx 101,332$ N. The atmospheric pressure at sea level is about 101,332 pascals, or 101.3 kilopascals (kPa).

That is a lot of pressure! Why does your rib cage not collapse? Because the air *inside* your lungs pushes out just as hard as the air outside pushes in. The two forces cancel, so you never notice them.

You can see the push when the balance breaks. Suck the air out of an empty plastic water bottle, and the atmosphere crushes it.

Exercise 4 The push on your hand

Your palm has an area of about 0.01 square meters (10 cm by 10 cm). What force does the atmosphere exert on it at sea level? What mass would weigh that much?

Working Space

Answer on Page 55

2.1 Altitude and Atmospheric Pressure

Let go of the balloon. As it rises, there is less and less air above it, so the atmospheric pressure on it drops. (The helium inside pushes the balloon bigger and bigger, until it pops.)




Figure 2.3: As a balloon rises, it reaches a point of less and less outer pressure.

What would be the atmospheric pressure at h meters above sea level? Here is a handy formula for that:

$$p = 101,332 \times \left(1 - \left(2.25577 \times 10^{-5} \times h\right)\right)^{5.25588}$$

where p is the atmospheric pressure in pascals.

Exercise 5 **Atmospheric Pressure**

Working Space

You are thinking about riding your bicycle to the top of Mount Everest. You worry that when the pressure outside the tire drops, the tire will burst. (A bursting tire is very, very loud.)

Calculate the atmospheric pressure at the top of Mount Everest (8,849 meters above sea level). How much does the outside pressure on the tire drop?

Answer on Page 56

2.2 How a Drinking Straw Works

When you suck on a straw, why does the drink rise? You might say you are pulling it up. Actually, the atmosphere is pushing it up.

Before you sip, the atmosphere presses evenly on the whole surface of the drink, inside the straw and out, so the surface stays level.

When you suck some air out of the straw, the pressure on the surface inside the straw drops. The atmospheric pressure on the surface outside the straw pushes into the straw and the beverage rises.





Figure 2.4: You use pressure to counteract gravity in a straw.

Gravity still pulls down on the liquid in the straw. The higher the liquid rises, the more liquid there is to hold up, and the harder you must suck.

2.2.1 The Longest Usable Straw

Assuming you are drinking water in a place with 100 kPa of atmospheric pressure, how high could you suck water with a perfect vacuum? That is, given a very, very long and very, very stiff drinking straw, if you created a pressure of 0 Pa inside, how far above the surface of the glass could you get the water?

Let us say a cross-section of the straw has an area of A square meters, and the very top of the column of water is h meters above the surface in the glass.



With how many newtons of force is the atmosphere pushing the water upward? $100 \text{ kPa} = 100,000 \text{ newtons per square meter}$. So:

$$F_u = (100,000)a$$

With how many newtons of force is gravity pulling the water in the straw downward? The volume of the water is ah . A cubic meter of water has a mass of 1000 kg . The force of gravity is $9.8 \text{ Newtons per kg}$.

$$F_d = (ah)(1000)(9.8)$$

The water will stop rising when $F_u = F_d$. So, to find h , we substitute in:

$$(100,000)a = (ah)(1000)(9.8)$$

Notice that we can divide both sides by ρg getting:

$$h = \frac{100,000}{9,800} = 10.2 \text{ meters}$$

Even a perfect vacuum could lift water only 10.2 meters, about three stories. That is why old-fashioned hand pumps cannot raise water from a well deeper than about 10 meters: they work by suction, and the atmosphere can only push so hard. Deep wells need a pump at the bottom that pushes the water up.


2.2.2 Millimeters Mercury

The density of mercury is 13,500 kg per cubic meter. How far up a straw would a perfect vacuum pull mercury?

$$h = \frac{100,000}{(9.8)(13,500)} \approx 0.756 \text{ meters}$$

When the atmospheric pressure is 100kPa, the mercury will rise 756 mm into a vacuum.

This is how a classic *barometer* measures atmospheric pressure. Fill a long glass tube with mercury, close one end, and stand it upside down, open end in a dish of mercury. The mercury drops until the atmosphere can hold it up. Read the height of the column off the millimeter marks. At sea level, it is about 760 mm, which is why weather reports sometimes give pressure in “millimeters of mercury.”



Atmospheric pressure rises and falls with the weather. Falling pressure often means a storm is on the way; rising pressure usually brings clear skies. A weather fan with a barometer at home might say, “The barometer went from 752 to 761 millimeters in the last hour. A high-pressure system is moving in.”

Exercise 6 A water barometer

Why do barometers use mercury instead of water? How tall would a water barometer’s column be at a sea-level pressure of 101,332 Pa?

Working Space

Answer on Page 56

2.3 How a Siphon Works

Put two cups on a table, one full of water and one empty, and connect them with an empty upside-down U of tubing. Nothing happens. The air pressure on both ends is the same, and climbing over the hump would take energy.

Now fill the tube with water first. Gravity pulls the water down on both sides of the


hump. If the outlet side is longer (it goes lower), it pulls harder. As that water falls, it leaves behind lower pressure at the top of the hump, and the atmosphere pushing on the full cup shoves more water up the short side to fill the gap.

The flow keeps itself going. Water runs from the full cup to the empty one until their surfaces are at the same level. Then the two sides balance, and the flow stops.

This is a *siphon*. Gravity drives it, but the atmosphere makes it possible. On a planet with gravity but no atmosphere, the water at the top of the hump would just break apart, and the siphon would fail.

Siphons are handy for emptying containers too big to tip, like a fish tank. Put one end of a hose in the tank, fill the hose with water (by dunking it all the way under, not by sucking on it), cover the end with your thumb, and lower that end into a bucket below the tank. Let go, and the tank drains into the bucket. (Never start a siphon by mouth with gasoline or any other chemical. Swallowing even a little gasoline can kill you.)






Siphons follow two rules:

- The peak of the siphon needs to be low enough for atmospheric pressure to push the liquid that high. For water at sea-level, for example, the peak of the siphon cannot be more than 10.2 meters above the surface of the source liquid.
- The tube has to carry the liquid to a lower level than the surface of the source liquid. If the destination end of the siphon is submerged, the surface of the liquid it is submerged in must be lower than the surface of the source liquid. If the destination end of the siphon is not submerged, its opening must be lower than the surface of the source liquid.

Follow those two rules, and you can get creative. In fact, every toilet has a siphon in it.

2.4 How a Toilet Works

First, a word about P-traps. Look under a sink, and you will see the drain pipe dip down and curve back up. This is a *P-trap*. Water stays trapped in the dip and blocks stinky (and flammable!) sewer gases from coming up into your house.



If a drain that has not been used in a while smells like sewage, the water in its trap has probably evaporated. Run some water to refill it.

Now, the toilet. The bowl drains through a curved tube that rises and then falls into the sewer: a siphon. Most of the time, that tube is partly full of air. When you flush, water rushes from the tank into the bowl and fills the tube. The siphon starts and sucks the bowl nearly empty, until air gets into the tube with a gurgle and the flow stops. That gurgle is the siphon breaking.




Figure 2.5: The steps a toilet takes when siphoning out waste.

The tank is simple. A float rides on the water. When the water level drops, the float drops and opens a valve that refills the tank. When the water is high enough, the float closes the valve. So the tank always holds the same amount of water, ready for the next flush.

When you flush, a flap in the bottom of the tank opens and the water rushes into the bowl. When the tank is empty, the flap closes and the tank refills.

Exercise 7 **Will it siphon?**

You want to siphon water from a pond over a wall into a field. The top of the wall is 12 meters above the pond's surface, and the field is 15 meters below the pond. Will the siphon work? Explain using the two rules.

Working Space

Answer on Page 56

Complex Numbers

In the last chapter, you met quadratics like $x^2 + 2x + 5 = 0$ that have no real solutions, because the formula asks for the square root of a negative number. For centuries, mathematicians simply gave up at that point. Then a few brave ones asked: what if we just *invent* a square root of -1 and see what happens? What happened was one of the most useful ideas in all of mathematics. Electrical engineers use it to design the power grid, and physicists use it to describe the atom.

3.1 Definition

The *imaginary unit* i is a number with the property

$$i^2 = -1$$

No real number does that, which is why i got the (unfortunate) name “imaginary.” It is no more imaginary than negative numbers, which also seemed absurd when they were new.

A *complex number* is a number of the form $a + bi$, where a and b are real numbers. The number a is the *real part*, and b is the *imaginary part*. For example, $3 - 2i$ has real part 3 and imaginary part -2 .

Every real number is a complex number with imaginary part 0: $7 = 7 + 0i$. A complex number with real part 0, like $5i$ or $-3i$, is called an *imaginary number*.

With i , every negative number has square roots. For example, $\sqrt{-9} = \sqrt{9}\sqrt{-1} = 3i$. Check: $(3i)^2 = 9i^2 = -9$.

Worked example. Solve $x^2 + 2x + 5 = 0$.


Solution. The quadratic formula gives

$$x = \frac{-2 \pm \sqrt{4 - 20}}{2} = \frac{-2 \pm \sqrt{-16}}{2} = \frac{-2 \pm 4i}{2} = -1 \pm 2i$$

Check $x = -1 + 2i$: $(-1 + 2i)^2 + 2(-1 + 2i) + 5 = (1 - 4i + 4i^2) - 2 + 4i + 5 = 1 - 4 - 2 + 5 = 0$. It works!

3.2 The Complex Plane

Real numbers live on a line. Complex numbers need a whole plane. Plot $a + bi$ at the point (a, b) : the real part goes left and right, and the imaginary part goes up and down.



The distance from 0 to $a + bi$ is the *magnitude* (or absolute value) of the number. By the Pythagorean theorem,

$$|a + bi| = \sqrt{a^2 + b^2}$$

For example, $|3 + 4i| = \sqrt{9 + 16} = 5$.

3.3 Adding and Subtracting

To add complex numbers, add the real parts and add the imaginary parts:

$$(a + bi) + (c + di) = (a + c) + (b + d)i \quad (3.1)$$

For example, $(2 + 3i) + (4 - i) = 6 + 2i$. Subtraction works the same way: $(2 + 3i) - (4 - i) = -2 + 4i$.

In the complex plane, adding is just adding vectors, tip to tail.

3.4 Multiplying

To multiply, expand just as you would $(x + 2)(x + 3)$, then replace i^2 with -1 :

$$(a + bi)(c + di) = ac + adi + bci + bdi^2 = (ac - bd) + (ad + bc)i \quad (3.2)$$

Worked example. Multiply $(2 + 3i)(1 - i)$.

Solution. $(2 + 3i)(1 - i) = 2 - 2i + 3i - 3i^2 = 2 + i + 3 = 5 + i$. The key step is $-3i^2 = -3(-1) = +3$.

Here is something beautiful. Multiplying by i *rotates* a number a quarter turn counter-clockwise around 0. Try it: $(2 + i) \cdot i = 2i + i^2 = -1 + 2i$. Plot both points, and you will see a 90° turn. Multiply by i twice, and you turn 180° , which is the same as multiplying by -1 . That is exactly why $i^2 = -1$.

It also explains the powers of i :

$$i^1 = i, \quad i^2 = -1, \quad i^3 = -i, \quad i^4 = 1, \quad i^5 = i, \dots$$

Every four quarter turns bring you back where you started.

Exercise 8 Square a Complex Number

Find $(1 + i)^2$. Then find $(1 + i)^4$.

Working Space

Answer on Page 56

3.5 Conjugates and Division

The *conjugate* of $a + bi$ is $a - bi$: same real part, opposite imaginary part. A number times its conjugate is always a real number:

$$(a + bi)(a - bi) = a^2 - abi + abi - b^2i^2 = a^2 + b^2$$

That is the trick for dividing. Multiply the top and bottom by the conjugate of the bottom, and the i disappears from the denominator.

Worked example. Compute $\frac{3+4i}{1+2i}$.

Solution. The conjugate of $1 + 2i$ is $1 - 2i$:

$$\frac{3 + 4i}{1 + 2i} \cdot \frac{1 - 2i}{1 - 2i} = \frac{3 - 6i + 4i - 8i^2}{1 + 4} = \frac{11 - 2i}{5} = 2.2 - 0.4i$$

Check by multiplying back: $(1 + 2i)(2.2 - 0.4i) = 2.2 - 0.4i + 4.4i - 0.8i^2 = 3 + 4i$.

3.6 Complex Numbers in Python

Python has complex numbers built in. Engineers write j instead of i (because i often stands for electric current), and Python follows them:

```

1 z = 2 + 3j
2 w = 1 - 1j
3 print(z * w)          # (5+1j)
4 print(abs(3 + 4j))    # 5.0
5 print(z.conjugate())  # (2-3j)

```

3.7 Why Complex Numbers Matter

Every polynomial has roots. The *fundamental theorem of algebra* says that every polynomial of degree n has exactly n complex roots (counting repeated roots more than once). With real numbers alone, even $x^2 + 1 = 0$ has no solution. With complex numbers, nothing is ever missing.

Rotations and waves. Since multiplying by a complex number rotates and stretches, complex numbers are a perfect language for anything that spins or oscillates. Electrical engineers describe alternating current with them, and the equations of quantum mechanics are written in them.

Exercise 9 Practice

Simplify each expression.

1. $(5 + 2i) - (3 - 4i)$
2. $(2 - i)(3 + 2i)$
3. i^{27}
4. $|-5 + 12i|$
5. $\frac{2+i}{1-i}$

Working Space

Answer on Page 56

Exercise 10 **Complex roots**

Solve $x^2 + 4x + 13 = 0$. Check one of your answers.

Working Space

Answer on Page 57

Exercise 11 **Quarter turns**

Plot $z = 3 + i$. Then plot iz , i^2z , and i^3z .
What shape do the four points make?

Working Space

Answer on Page 57

Introduction to Sequences

Which would you rather have: a million dollars today, or a penny today, two pennies tomorrow, four the day after, doubling every day for 30 days? Take the pennies. On day 30 alone, you would get more than five million dollars. By the end of this chapter, you will be able to prove it.

A *sequence* is a list of numbers in a particular order. $\{1, 3, 5, 7, 9\}$ is a sequence. So is $\{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots\}$, and so are your daily pennies: $\{1, 2, 4, 8, \dots\}$. This chapter covers the two most common kinds: arithmetic and geometric sequences.

Sequences are generally represented like this:

$$a_1, a_2, a_3, a_4, \dots, a_n, \dots$$

The first number, a_1 , is called the *first term*, a_2 is the *second term*, and a_n is the *nth term*. A sequence can be finite or infinite. If the sequence is infinite, we represent that with ellipses ($\dots$) at the end of the list, to indicate that there are more numbers.

We can also write formulas to represent a sequence. Take the first example, the finite sequence $\{1, 3, 5, 7, 9\}$. Notice that each term is two more than the previous term. We can define the sequence *recursively* by defining the n^{th} term as a function of the $(n-1)^{\text{th}}$ term. In our example, we see that $a_n = a_{n-1} + 2$ with $a_1 = 1$ for $1 \leq n \leq 5$. This is called a recursive formula, because you have to already know the $(n-1)^{\text{th}}$ term to find the n^{th} term.

Another way to write a formula for a sequence is to find a rule for the n^{th} term. In our example sequence, the first term is 1 plus 0 times 2, the second term is 1 plus 1 times 2, the third term is 1 plus 2 times 2, and so on. Did you notice the pattern? The n^{th} term is 1 plus $(n-1)$ times 2. We can write this mathematically:

$$a_n = 1 + 2(n-1) \text{ for } 1 \leq n \leq 5$$

This is called the *explicit* formula because each term is explicitly defined. Notice that for the second way of writing a formula, we do not have to state what the first term is — the formula tells us.

4.1 Arithmetic sequences

Our first example sequence, $\{1, 3, 5, 7, 9\}$ is a *finite, arithmetic* sequence. We know it is finite because there is a limited number of terms in the sequence (in this case, 5). How do we know it is arithmetic?

An arithmetic sequence is one where you add the same number every time to get the next term. Our example is an arithmetic sequence because you add 2 to get the next term every time. That number that you add is called the *common difference*, so we can say the sequence $\{1, 3, 5, 7, 9\}$ has a common difference of 2. The common difference can be positive (in the case of an increasing arithmetic sequence) or negative (in the case of a decreasing arithmetic sequence). Formally, we can find the common difference of an arithmetic sequence by subtracting the $(n - 1)^{\text{th}}$ term from the n^{th} term:

$$d = a_n - a_{n-1}$$

Exercise 12

Which of the following are arithmetic sequences? For the arithmetic sequences, find the common difference.

1. $\{\frac{1}{2}, \frac{1}{4}, \frac{1}{6}, \frac{1}{8}, \dots\}$
2. $\{5, 8, 11, 14, 17, \dots\}$
3. $\{3, -1, -5, -9, \dots\}$
4. $\{-1, 2, -3, 4, -5, 6, \dots\}$

Working Space

Answer on Page 57

4.1.1 Formulas for arithmetic sequences

If you are given an arithmetic sequence, you can write an explicit or recursive formula. You can think of the formula as a function where the domain (input) is restricted to integers greater than or equal to one. Let us write explicit and recursive formulas for the sequence $\{3, -1, -5, -9, \dots\}$.

For either type of formula, we need to identify the common difference. Since each term

is 4 less than the previous term, the common difference is -4 (see figure 4.1). This means the n^{th} term is the $(n - 1)^{\text{th}}$ term minus 4. The general form of a recursive formula is $a_n = a_{n-1} + d$, where d is the common difference. For our example, the common difference is -4, so we can write a recursive formula:

$$a_n = a_{n-1} - 4$$

However, this formula does not tell us what a_1 is! For recursive formulas, you have to specify the first term in the sequence. So, the *complete* recursive formula for the sequence is:

$$\begin{aligned} a_n &= a_{n-1} - 4 \\ a_1 &= 3 \end{aligned}$$




Figure 4.1: The common difference in the sequence $\{3, -1, -5, -9, \dots\}$ is -4

Recursive formulas make it easy to see how each term is related to the next term. However, it is difficult to use recursive formulas to find a specific term. Say we wanted to know the 7th term in the sequence. Well, from the formula, we know that:

$$a_7 = a_6 - 4$$

What is a_6 ? Again, we see that

$$a_6 = a_5 - 4$$

Now we have to find a_5 ! If we keep going, we see that:

$$a_5 = a_4 - 4$$

$$a_4 = a_3 - 4$$

$$a_3 = a_2 - 4$$

$$a_2 = a_1 - 4$$

Since we were told a_1 , we can find a_2 and propagate our terms back up the chain to find a_7 :

$$a_2 = 3 - 4 = -1$$

$$a_3 = a_2 - 4 = -1 - 4 = -5$$

$$a_4 = a_3 - 4 = -5 - 4 = -9$$

$$a_5 = a_4 - 4 = -9 - 4 = -13$$

$$a_6 = a_5 - 4 = -13 - 4 = -17$$

$$a_7 = a_6 - 4 = -17 - 4 = -21$$

Ultimately, we see that $a_7 = -21$. That was a lot of work! You can imagine that for higher n terms, such as the 100^{th} or 1000^{th} term, this method becomes cumbersome. This is where the explicit formula is more useful.

The general form of an explicit formula for an arithmetic sequence is

$$a_n = a_1 + d \times (n - 1)$$

where d is the common difference. For our example sequence, $\{3, -1, -5, -9, \dots\}$, the common difference is -4 . So the explicit formula is

$$a_n = 3 + (-4)(n - 1) = 3 - 4(n - 1)$$

You may be tempted to distribute and simplify, which is fine and yields an equivalent formula:

$$a_n = 7 - 4n$$

Now, to find the 7^{th} term, all we have to do is substitute $n = 7$:

$$a_7 = 3 - 4(7 - 1) = 3 - 4(6) = 3 - 24 = -21$$

We get the same answer with much less effort!

Worked example. Movie tickets cost \$12 each, and the theater charges a \$3 online booking fee per order. Write a formula for the cost of n tickets, and find the cost of 10.

Solution. The costs for 1, 2, 3, $\dots$ tickets are $\{15, 27, 39, \dots\}$: an arithmetic sequence with $a_1 = 15$ and $d = 12$. So $a_n = 15 + 12(n - 1) = 3 + 12n$, and $a_{10} = 3 + 120 = \$123$.

Not every recursive formula makes an arithmetic sequence. The famous Fibonacci sequence, which you met in the Python chapters, adds the two previous terms: $a_n = a_{n-1} + a_{n-2}$, with $a_1 = a_2 = 1$. That gives $\{1, 1, 2, 3, 5, 8, 13, \dots\}$. The differences keep changing, so it is not arithmetic.

Exercise 13

An arithmetic sequence is defined by the recursive formula $a_n = a_{n-1} + 5$ with $a_1 = -4$. Write the first 5 terms of the sequence and determine an explicit formula for the same sequence.

Working Space

Answer on Page 57

Exercise 14

The first four terms of an arithmetic sequence are $\{\pi, \frac{3\pi}{2}, 2\pi, \frac{5\pi}{2}\}$. What is the common difference? Write explicit and recursive formulas for the infinite sequence.

Working Space

Answer on Page 57

4.2 Geometric sequences

Let us look at the other sequence given as an example at the beginning of the chapter: $\{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots\}$. How is each term related to the previous term? Well, $\frac{1}{4}$ is half of $\frac{1}{2}$, and $\frac{1}{8}$ is half of $\frac{1}{4}$, so each term is the previous term multiplied by $\frac{1}{2}$. When you get each term by multiplying the previous term by the same number, you have a *geometric* sequence. The number you multiply by (here, $\frac{1}{2}$) is called the *common ratio*. The common ratio can be positive or negative, but not zero.

An easy way to determine the common ratio (r) is to divide the n^{th} term by the $(n-1)^{\text{th}}$ term. In our example sequence, the first term is $\frac{1}{2}$ and the second is $\frac{1}{4}$.

$$r = \frac{a_2}{a_1} = \frac{1/4}{1/2} = \frac{1}{2}$$

which returns the common ratio we already identified, $r = \frac{1}{2}$.

If the common ratio is negative, then the sequence will “flip” back and forth from positive to negative. For example, suppose there is a geometric sequence such that $a_1 = 1$ and $r = -2$. Then the first 5 terms are $\{1, -2, 4, -8, 16\}$. Whenever you see a sequence going back and forth from positive to negative, that means the common ratio is negative.

If the first term is positive and $r > 1$, the terms grow, faster and faster. If $0 < r < 1$, they shrink toward zero. If r is negative, they flip signs every step.

4.2.1 Formulas for geometric sequences

Like arithmetic sequences, we can write recursive and explicit formulas. For geometric sequences, the recursive formula has the general form:

$$a_n = r(a_{n-1})$$

where r is the common ratio and a_1 is specified. For our example sequence, $\{\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \dots\}$, the recursive formula is:

$$\begin{aligned}a_n &= \frac{1}{2}a_{n-1} \\ a_1 &= \frac{1}{2}\end{aligned}$$

In a geometric sequence, each term is the first term, a_1 , multiplied by the common ratio, r , $n - 1$ times. Therefore, the general form of an explicit formula for a geometric sequence is:

$$a_n = (a_1)r^{n-1}$$

Again, for our example sequence, $a_1 = \frac{1}{2}$ and $r = \frac{1}{2}$, so the explicit formula is:

$$a_n = \left(\frac{1}{2}\right)\left(\frac{1}{2}\right)^{(n-1)}$$

Exercise 15

Which of the following are geometric sequences? For each geometric sequence, determine the common ratio.

Working Space

1. $\{2, -4, 6, -8, \dots\}$

2. $\{4, 2, 1, \frac{1}{2}, \dots\}$
3. $\{-5, 25, -125, 625, \dots\}$
4. $\{2, 0, -2, -4, \dots\}$

Answer on Page 58

Exercise 16

A geometric sequence is defined by the recursive formula $a_n = a_{n-1} \times \frac{3}{2}$ with $a_1 = 1$. Write the first five terms of the sequence and determine an explicit formula for the same sequence.

Working Space

Answer on Page 58

Exercise 17

The first four terms of a geometric sequence are $\{-4, 2, -1, \frac{1}{2}\}$. What is the common ratio? Write recursive and explicit formulas for the infinite sequence.

Working Space

Answer on Page 58

Exercise 18 Pennies from heaven

Back to the opening question. On day 1 you get \$0.01, and each day you get twice as much as the day before. Write an explicit formula for the amount you

Working Space

get on day n . How much do you get on day 30?

Answer on Page 58

Exercise 19 **Bouncing ball**

A ball is dropped from 2 meters. Each bounce reaches 80% of the height of the one before. Write an explicit formula for the height of the n th bounce. How high is the 5th bounce? Is this sequence arithmetic, geometric, or neither?

Working Space

Answer on Page 58

Exponentials and Logarithms

Fold a sheet of paper in half, and it is twice as thick. Fold it again: four times. If you could fold it 42 times, it would reach the Moon. That is the power of repeated multiplication, and this chapter is about the tools for handling it: exponents, and their opposites, logarithms. With them, you can figure out when your savings will double, how long a glowing watch keeps glowing, and how fast a video is going viral.

5.1 Exponents

Let us start with a quick review. Long ago, scientists noticed that their formulas kept multiplying the same number by itself. For example, if they knew that a sphere was r centimeters in radius, its volume in milliliters was

$$V = \frac{4}{3} \times \pi \times r \times r \times r$$

They did two things to make the notation less messy. First, they decided that if two numbers were written next to each other, the reader would assume that meant “multiply them”. Second, they came up with the exponent, a little number that was lifted off the baseline of the text, that meant “multiply it by itself”. For example 5^3 was the same as $5 \times 5 \times 5$.

Now, the formula for the volume of a sphere is written

$$V = \frac{4}{3}\pi r^3$$

Tidy, right? In an exponent expression like this, we say that r is *the base* and 3 is *the exponent*.

5.1.1 Identities for Exponents

What about exponents of exponents? What is $(5^3)^2$?

$$(5^3)^2 = (5 \times 5 \times 5)^2 = (5 \times 5 \times 5)(5 \times 5 \times 5) = 5^6$$

In general, for any a , b , and c :

$$(a^b)^c = a^{(bc)}$$

If you have $(5^3)(5^4)$, that is just $5 \times 5 \times 5 \times 5 \times 5 \times 5 \times 5$ or 5^7

The general rule is that for any a , b , and c

$$(a^b)(a^c) = a^{(b+c)}$$

Mathematicians *love* this rule, so we keep extending the idea of exponents to keep this rule true. For example, at some point, someone asked “What about 5^0 ?” According to the rule, 5^2 must equal $5^{(2+0)}$ which must equal $(5^2)(5^0)$. Thus, 5^0 must be 1. So mathematicians declared “Anything to the power of 0 is 1”.

We do not typically assume that $0^0 = 1$. It is just too weird. So, we say that for any a not equal to zero,

$$a^0 = 1$$

What about $5^{(-2)}$? By our beloved rule, we know that $(5^{-2})(5^5)$ must be equal to 5^3 , right? So 5^{-2} must be equal to $\frac{1}{5^2}$.

We say that for any a not equal to zero and any b ,

$$a^{-b} = \frac{1}{a^b}$$

This makes dividing one exponential expression by another (with the same base) easy:

$$\frac{a^b}{a^c} = a^{(b-c)}$$

We often say “cancel out” for this. Here, we can “cancel out” x^2 :

$$\frac{x^5}{x^2} = x^3$$

What about $5^{\frac{1}{3}}$? By the beloved rule, we know that $5^{\frac{1}{3}}5^{\frac{1}{3}}5^{\frac{1}{3}}$ must equal 5^1 . Thus, $5^{\frac{1}{3}} = \sqrt[3]{5}$.

We say that for any positive a , any b , and any positive integer c ,

$$a^{\frac{b}{c}} = \left(\sqrt[c]{a}\right)^b = \left(\sqrt[b]{a^c}\right)$$

Before you go on to the exercises, note that the beloved rule demands a common base.

- We can combine these: $(5^2)(5^4) = 5^6$
- We cannot combine: $(5^2)(3^5)$

With that said, we note for any a, b , and c :

$$(ab)^c = (a^c)(b^c)$$

So, for example, if we were asked to simplify $(3^4)(6^2)$, we would note that $6 = 2 \times 3$, so

$$(3^4)(6^2) = (3^4)(3^2)(2^2) = (3^6)(2^2)$$

If these ideas are new to you (or maybe you are just having trouble remembering them), watch the Khan Academy's **Intro to rational exponents** video at <https://youtu.be/1ZfXc4nHooo>.

5.2 Logarithms

After the world created exponents, it needed the opposite. We could talk about the quantity $? = 2^3$, that is, "What is the product of 2 multiplied by itself three times?" We needed some way to talk about $2^? = 8$, that is, "2 to the what is 8?" This is why we developed the logarithm.

Here is an example:

$$\log_2 8 = 3$$

In English, you would say, “The logarithm base 2 of 8 is 3.”

The base (2, in this case) can be any positive number. The argument (8, in this case) can also be any positive number.

Try this one: What is the logarithm base 2 of $1/16$?

You know that $2^{-4} = \frac{1}{16}$, so $\log_2 \frac{1}{16} = -4$.

Worked example. Find $\log_2 1024$ and $\log_{10} 0.001$.

Solution. $2^{10} = 1024$, so $\log_2 1024 = 10$. And $10^{-3} = 0.001$, so $\log_{10} 0.001 = -3$. A logarithm just answers the question “what power?”

Exercise 20 Solve for the Exponent

Solve $3^x = 81$. Then find $\log_5 125$ and $\log_4 2$.

Working Space

Answer on Page 58

5.2.1 Logarithms in Python

Most calculators have limited logarithm buttons, but Python has a nice `log` function that lets you specify both the argument and the base. Start python, import the `math` module, and try taking a few logarithms:

```
1 >>> import math
2 >>> math.log(8,2)
3 3.0
4 >>> math.log(1/16, 2)
5 -4.0
```

Say a friend offers you 5% interest per year on your investment for as long as you want. Each year, your money is multiplied by 1.05. You wonder, “How many years before my investment is 100 times as large?” You can solve this problem with logarithms:

```
1 >>> math.log(100, 1.05)
```

2 | 94.3872656381287

If you leave your investment with your friend for 94.4 years, the investment will be worth 100 times what you put in.

5.2.2 Logarithm Identities

The logarithm is defined this way:

$$\log_b a = c \iff b^c = a$$

Notice that the logarithm of 1 is always zero, and $\log_b b = 1$.

The logarithm of a product:

$$\log_b ac = \log_b a + \log_b c$$

This follows from the exponent rule $b^{x+y} = b^x b^y$: logarithms are exponents, so multiplying numbers adds their logarithms. What about a quotient?

$$\log_b \frac{a}{c} = \log_b a - \log_b c$$

Exponents?

$$\log_b (a^c) = c \log_b a$$

Notice that because logs and exponents are the opposite of each other, they can cancel each other out:

$$b^{\log_b a} = a$$

and

$$\log_b (b^a) = a$$

5.2.3 Changing Bases

We mentioned that most calculators have pretty limited logarithm capabilities. Most calculators do not allow you to specify what base you want to work with. All scientific calculators have a button for “log base 10”. So, you need to know how to use that button to get logarithms for other bases. Here is the change-of-base identity:

$$\log_b a = \frac{\log_c a}{\log_c b}$$

For example, if you wanted to find $\log_2 8$, you would ask the calculator for $\log_{10} 8$, then divide that by $\log_{10} 2$. You should get 3.

5.2.4 Natural Logarithm

The circumference of a circle is about 3.141592653589793 times its diameter. Because we use this unwieldy number a lot, we give it a name: We say “The circumference of a circle is π times its diameter.”

There is a second unwieldy number that we will eventually use a great deal in solving problems. It is about 2.718281828459045 (but the digits actually go on forever, just like π). We call this number e . (You will see where e comes from in the section on continuous compound interest.)

Most calculators have a button labeled “ln”. That is the *natural logarithm* button. It takes the log in base e .

Similarly, in python, if you do not specify a base, the logarithm is done in base e :

```
1 >>> math.log(10)
2 2.302585092994046
3 >>> math.log(math.e)
4 1.0
```

5.2.5 Logarithms in Spreadsheets

Spreadsheets have three log functions:

- LOG takes both the argument and the base. LOG(8,2) returns 3.

- LOG10 takes just the argument and uses 10 as the base.
- LN takes just the argument and uses e as the base.

Here is a plot from a spreadsheet of a graph of $y = \text{LOG}(x, 2)$.




Figure 5.1: A graph of log base 2.

Spreadsheets also have the function $\text{EXP}(x)$, which returns e^x . For example, $\text{EXP}(2)$ returns 7.38905609893065.

5.3 Compound Interest

When you lend money, you usually charge interest. You may not think of it this way, but a savings account is a loan: you lend your money to the bank. Any money you put in the account is loaned to the bank. The bank then lends it to someone else, who pays interest to the bank. The bank gives some of that interest to you. However, what if you leave the interest in your account? And you start making *interest on the interest*? This is known as *compound interest*.

5.3.1 An example with annual interest payments

Let us say that you put \$1000 in a savings account that pays 6% interest every year. How much money would you have after 12 years? Let us make a spreadsheet.

	A	B	C
1	Interest Rate	6.00%	
2			
3	After year:	Interest	Balance
4	0	\$0.00	1000
5	1	\$60.00	\$1,060.00

Create a new spreadsheet and edit the cells to look like this. All the cells in rows 1 - 4 are just values: just type in what you see.

The fifth row is all formulas:

After year	Interest	Balance
= A4 + 1	= B\$1 * C4	= C4 + B5

The interest rate field should be formatted as a percentage. One thing to know when dealing with percentages in the spreadsheet: If the field says "600%", its value is 6.

The cells in the Interest and Balance column should be formatted as currency.

You are about to make several copies of the cells in the fifth row, so make sure they look right.

Click on A5 and shift-click on C5 to select all three cells. Drag the lower-right corner down to fill the rows 6 - 15.

A5:C16 fx = A4 + 1			
	A	B	C
1	Interest Rate	6.00%	
2			
3	After year	Interest	Balance
4	0	\$0.00	1000
5	1	\$60.00	\$1,060.00
6	2	\$63.60	\$1,123.60
7	3	\$67.42	\$1,191.02
8	4	\$71.46	\$1,262.48
9	5	\$75.75	\$1,338.23
10	6	\$80.29	\$1,418.52
11	7	\$85.11	\$1,503.63
12	8	\$90.22	\$1,593.85
13	9	\$95.63	\$1,689.48
14	10	\$101.37	\$1,790.85
15	11	\$107.45	\$1,898.30
16	12	\$113.90	\$2,012.20
17			

Look at the numbers. The first interest payment is \$60, but the last is \$113.90. Your balance has more than doubled!

5.3.2 Exponential Growth

We figured this out numerically by repeatedly multiplying the balance by the interest rate. What if you wanted to know what the balance would be n years after investing P_0 dollars with an annual interest rate of r ? (Note that r in our example would be 0.06, not 6.0.)

Each year, the balance is multiplied by $1+r$, so after one year, P_0 would become $P_0 \times (1+r)$. The next year you would multiply this number by $(1+r)$ again: $P_0 \times (1+r) \times (1+r)$. The next year? $P_0 \times (1+r) \times (1+r) \times (1+r)$ See the pattern? P_n is this balance after n years, then

$$P_n = P_0(1+r)^n \quad (5.1)$$

5.1 is referred to as the standard *compound interest formula*. Because n is an exponent, we call this *exponential growth*. There are few things as terrifying to a scientist as the phrase “The population is undergoing exponential growth”.

Many banks compound more often than once a year. If the bank compounds k times a year (12 for monthly, 365 for daily), it pays $\frac{r}{k}$ each time, and it does that kt times in t years. Your principal P_0 grows to

$$P = P_0 \left(1 + \frac{r}{k}\right)^{kt}$$

5.3.3 Sensitivity to interest rate

For most people, the first surprising thing about compound interest is how quickly your money grows after a few years. The second thing that is surprising is how much difference a small change in the percentage rate makes.

Let us add another set of columns that shows what happens to your money if you convince the bank to pay you 8% instead of 6%.

Copy everything from columns B and C:

	A	B	C	D	E
1	Interest Rate	6.00%		6.00%	
2					
3	After year	Interest	Balance	Interest	Balance
4	0	\$0.00	1000	\$0.00	1000
5	1	\$60.00	\$1,060.00	\$60.00	\$1,060.00
6	2	\$63.60	\$1,123.60	\$63.60	\$1,123.60
7	3	\$67.42	\$1,191.02	\$67.42	\$1,191.02
8	4	\$71.46	\$1,262.48	\$71.46	\$1,262.48
9	5	\$75.75	\$1,338.23	\$75.75	\$1,338.23
10	6	\$80.29	\$1,418.52	\$80.29	\$1,418.52
11	7	\$85.11	\$1,503.63	\$85.11	\$1,503.63
12	8	\$90.22	\$1,593.85	\$90.22	\$1,593.85
13	9	\$95.63	\$1,689.48	\$95.63	\$1,689.48
14	10	\$101.37	\$1,790.85	\$101.37	\$1,790.85
15	11	\$107.45	\$1,898.30	\$107.45	\$1,898.30
16	12	\$113.90	\$2,012.20	\$113.90	\$2,012.20
17					

Now edit the second interest rate to be 8%:

	A	B	C	D	E
1	Interest Rate	6.00%		8.00%	
2					
3	After year	Interest	Balance	Interest	Balance
4	0	\$0.00	1000	\$0.00	1000
5	1	\$60.00	\$1,060.00	\$80.00	\$1,080.00
6	2	\$63.60	\$1,123.60	\$86.40	\$1,166.40
7	3	\$67.42	\$1,191.02	\$93.31	\$1,259.71
8	4	\$71.46	\$1,262.48	\$100.78	\$1,360.49
9	5	\$75.75	\$1,338.23	\$108.84	\$1,469.33
10	6	\$80.29	\$1,418.52	\$117.55	\$1,586.87
11	7	\$85.11	\$1,503.63	\$126.95	\$1,713.82
12	8	\$90.22	\$1,593.85	\$137.11	\$1,850.93
13	9	\$95.63	\$1,689.48	\$148.07	\$1,999.00
14	10	\$101.37	\$1,790.85	\$159.92	\$2,158.92
15	11	\$107.45	\$1,898.30	\$172.71	\$2,331.64
16	12	\$113.90	\$2,012.20	\$186.53	\$2,518.17
17					

5.3.4 Continuous Compound Interest

Monthly beats yearly, and daily beats monthly. What if the bank compounded every hour? Every second? Does your money grow without limit?

Try it with a 100% interest rate for one year, starting with \$1. If the bank compounds k times, you end with $(1 + \frac{1}{k})^k$ dollars:

Compounding	k	$(1 + \frac{1}{k})^k$
yearly	1	2
monthly	12	2.613...
daily	365	2.7146...
every second	31,536,000	2.7182818...

The answers creep up, but they never pass 2.71828.... That number is e ! It is the most money \$1 can earn in a year at 100% interest, no matter how often the bank compounds. This is the real reason e shows up everywhere growth and decay happen.

With interest rate r for t years, compounding “infinitely often” gives the *continuous compound interest* formula:

$$P = P_0 e^{rt} \quad (5.2)$$

(In the calculus volumes, you will learn to make “infinitely often” precise with a limit.)

5.3.5 Practice Exercises

Exercise 21

You deposit \$1500 into a savings account that pays an annual interest rate of 5%, compounded annually. What is the balance after 8 years?

Note: you can also use excel or google sheets built in function =FV().

Working Space

Answer on Page 59

Exercise 22

A bank offers two savings accounts, both with a stated annual interest rate of 6%. Account A compounds monthly ($k = 12$) and Account B compounds daily ($k = 365$). If you deposit \$3000 into each account, which account has the higher balance after 5 years, and by how much?

Create two columns in Excel/Sheets showing the balance after each year for each account, using \$3000 as the principal. Start with year 0 for the initial deposit, then fill in years 1 through 5 using the compound interest formula for monthly and daily compounding.

Working Space

Answer on Page 59

5.4 Exponential Decay

In the last section, an investment of P earning compound interest at an annual rate r grew exponentially. At the end of year t , the balance is

$$P(1 + r)^t$$

Because r is positive, this number grows as time passes. You get a nice exponential growth curve that looks something like this:




Figure 5.2: A diagram showing (exponential) compound interest.

This is \$30 invested with a 10% annual interest rate. So, the formula for the balance after t years would be

$$(30)(1.1)^t$$

What if r were negative? This would be *exponential decay*.

5.4.1 Radioactive Decay

Until around 1970, there were companies making watches whose faces and hands were coated with radioactive paint. The paint usually contained radium. When a radium atom decays, it gives off some energy, loses two protons and two neutrons, and becomes a different element (radon). Some of the energy given off is visible light. Thus, these watches glow in the dark.

How many of the radium atoms in the paint decay each century? About 4.24%.

Notice the quantity of atoms lost is proportional to the number of atoms you have. This is exponential decay. If we assume that we start with a million radium atoms, the number of atoms decreases over time like this:




Figure 5.3: A diagram showing the decay and half life of radium.

- We start with 1,000,000 atoms.
- At 16 centuries, we have only 500,000 (half as many) left.
- 16 centuries after that, we have only 250,000 (half again) left.
- 16 centuries after that, we have only 125,000 (half again) left.

A nuclear chemist would say that radium has a *half-life* of 1,600 years: every 1,600 years, half of whatever radium remains decays. (Check: $(1 - 0.0424)^{16} \approx 0.50$.) This means that if you bought a watch with glowing hands in 1960, it will be glowing half as brightly in the year 3560.

How do we calculate the amount of radium left at the end of century t ? If you start with P atoms, at the end of the t -th century, you will have

$$P(1 - 0.0424)^t$$

This is exponential decay.

5.4.2 Model Exponential Decay

Exercise 23 A shrinking company

You are hired to run a company with 480,000 employees. Each year, $\frac{1}{8}$ of the employees leave (they retire, quit, or move away). For some reason, you never hire anyone new.

Make a spreadsheet showing how many of the original 480,000 employees are still around at the end of each of the next 12 years, and make a bar graph of the data. Then write a formula for the number left after t years, and use logarithms to find when the company first drops below half its size.

Working Space

Answer on Page 60

Real data often curves. Populations grow exponentially, and radioactive substances decay exponentially. You can force a line of best fit through curved data, but it will be the wrong model. Logarithms can straighten curves out, so you can use the line-fitting tools you already know.

Exponential growth and decay show up in data, too: video views, bacteria counts, and medicine in the bloodstream. After you learn to fit straight lines to data, the Transforming Nonlinear Data with Logarithms chapter shows how logarithms turn curves like these into straight lines that you can fit.

Answers to Exercises

Answer to Exercise 1 (on page 4)

```
1 print(scores["score"].max())
```

The output is 92.

Answer to Exercise 2 (on page 10)

```
1 big = df['size'] >= 57
2 print(big.sum())
```

It prints 313.

Answer to Exercise 3 (on page 10)

```
1 best = df['purchase_price'].idxmax()
2 print(best, df.loc[best, 'brand'])
```

It prints 5728868 Canyon: a \$377.70 Canyon.

Answer to Exercise 4 (on page 15)

$F = pA = (101,332)(0.01) \approx 1,013$ N. A mass of $\frac{1,013}{9.8} \approx 103$ kg weighs that much. The air pushes on the back of your hand just as hard, so the forces cancel.

Answer to Exercise 5 (on page 16)

$$p = 101,332 \times \left(1 - 2.25577 \times 10^{-5} \times h\right)^{5.25588}$$

and $h = 8,849$. Thus,

$$p \approx 31.4 \text{ kPa}$$

That is less than a third of the pressure at sea level. The pressure outside the tire drops by about $101.3 - 31.4 \approx 70$ kPa, so the tire pushes outward 70 kPa harder than before. A road bike tire is pumped to about 600 kPa above the outside air, so this is a big change, but most tires can take it. Climbers' bigger worry is breathing: each breath holds less than a third as much oxygen.

Answer to Exercise 6 (on page 20)

$h = \frac{101,332}{(1000)(9.8)} \approx 10.3$ meters. A water barometer would be taller than a three-story building. Mercury is 13.5 times as dense, so its column is 13.5 times shorter, small enough to fit on a wall.

Answer to Exercise 7 (on page 24)

No. The outlet is lower than the pond, so the second rule is satisfied. But the hump is 12 meters above the pond's surface, and the atmosphere can push water up only about 10.2 meters. The water column will break at the top, and the siphon will not run.

Answer to Exercise 8 (on page 27)

$$(1 + i)^2 = 1 + 2i + i^2 = 1 + 2i - 1 = 2i. \text{ So } (1 + i)^4 = (2i)^2 = 4i^2 = -4.$$

Answer to Exercise 9 (on page 28)

1. $2 + 6i$

2. $6 + 4i - 3i - 2i^2 = 8 + i$
3. $27 = 4 \cdot 6 + 3$, so $i^{27} = i^3 = -i$.
4. $\sqrt{25 + 144} = 13$
5. $\frac{(2+i)(1+i)}{(1-i)(1+i)} = \frac{2+3i+i^2}{2} = \frac{1+3i}{2} = 0.5 + 1.5i$

Answer to Exercise 10 (on page 29)

$$x = \frac{-4 \pm \sqrt{16-52}}{2} = \frac{-4 \pm \sqrt{-36}}{2} = \frac{-4 \pm 6i}{2} = -2 \pm 3i.$$

Check $-2 + 3i$: $(-2 + 3i)^2 = 4 - 12i + 9i^2 = -5 - 12i$, and $4(-2 + 3i) = -8 + 12i$. Adding 13: $-5 - 12i - 8 + 12i + 13 = 0$.

Answer to Exercise 11 (on page 29)

$iz = 3i + i^2 = -1 + 3i$, $i^2z = -3 - i$, and $i^3z = 1 - 3i$. Each is a quarter turn from the last around 0, so the four points are the corners of a square centered at 0.

Answer to Exercise 12 (on page 32)

1. not arithmetic
2. arithmetic, common difference is 3
3. arithmetic, common difference is -4
4. not arithmetic

Answer to Exercise 13 (on page 35)

The first five terms are $\{-4, 1, 6, 11, 16\}$ and an explicit formula is $a_n = -4 + 5(n - 1)$.

Answer to Exercise 14 (on page 35)

The common difference is $\frac{3\pi}{2} - \pi = \frac{\pi}{2}$. The recursive formula is $a_n = a_{n-1} + \frac{\pi}{2}$ with $a_1 = \pi$. The explicit formula is $a_n = \pi + \frac{\pi}{2}(n - 1)$.

Answer to Exercise 15 (on page 36)

1. not geometric ($\frac{-4}{2} = -2$, but $\frac{6}{-4} = -1.5$)
2. geometric sequence with common ratio $r = \frac{1}{2}$
3. geometric sequence with common ratio $r = -5$
4. not geometric

Answer to Exercise 16 (on page 37)

The first five terms are $\{1, \frac{3}{2}, \frac{9}{4}, \frac{27}{8}, \frac{81}{16}\}$. An explicit formula for this sequence is $a_n = 1(\frac{3}{2})^{(n-1)}$.

Answer to Exercise 17 (on page 37)

The common ratio is $\frac{a_n}{a_{n-1}} = \frac{2}{-4} = -\frac{1}{2}$. A recursive formula would be $a_n = a_{n-1} \times -\frac{1}{2}$ with $a_1 = -4$. An explicit formula would be $a_n = (-4)(-\frac{1}{2})^{(n-1)}$.

Answer to Exercise 18 (on page 37)

This is geometric with $a_1 = 0.01$ and $r = 2$, so $a_n = 0.01 \cdot 2^{n-1}$. On day 30, you get $0.01 \cdot 2^{29} = \$5,368,709.12$. (Add up all 30 days, and the total is more than \$10.7 million.)

Answer to Exercise 19 (on page 38)

The first bounce reaches $2(0.8) = 1.6$ m, and each bounce is 0.8 times the one before, so the sequence is geometric with $a_1 = 1.6$ and $r = 0.8$: $a_n = 1.6(0.8)^{n-1}$. The 5th bounce reaches $1.6(0.8)^4 \approx 0.66$ m.

Answer to Exercise 20 (on page 42)

$81 = 3^4$, so $x = 4$. Since $5^3 = 125$, $\log_5 125 = 3$. Since $4^{1/2} = 2$, $\log_4 2 = \frac{1}{2}$.

Answer to Exercise 21 (on page 49)

Using 5.1 with $P_0 = 1500$, $r = 0.05$, and $n = 8$:

$$P_8 = 1500(1 + 0.05)^8 = 1500(1.05)^8 \approx \$2216.18$$

Alternatively, in Sheets/Excel, we can use `=FV(0.05, 8, 0, -1500)` to say that we want the *future value* of 1500 deposited (hence, negative), after 8 years compounded once per year.

Answer to Exercise 22 (on page 50)

Using $P = P_0(1 + \frac{r}{k})^{kt}$ with $P_0 = 3000$, $r = 0.06$, and $t = 5$ years:

Account A (monthly, $k = 12$):

$$P = 3000 \left(1 + \frac{0.06}{12}\right)^{12 \times 5} \approx \$4046.55$$

Account B (daily, $k = 365$):

$$P = 3000 \left(1 + \frac{0.06}{365}\right)^{365 \times 5} \approx \$4049.48$$

	A	B	C	D
1	Principal	\$ 3,000.00		
2	Annual Rate	6%		
3				
4	Year	Account A	Account B	
5	0 - initial	\$3,000	\$3,000	
6	1	\$3,185.03	\$ 3,185.49	
7	2	\$3,381.48	\$ 3,382.46	
8	3	\$3,590.04	\$ 3,591.60	
9	4	\$3,811.47	\$ 3,813.67	
10	5	\$4,046.55	\$ 4,049.48	

Account B has the higher balance, by about \$2.93. This makes sense: the more frequently interest is compounded, the closer the balance gets to the continuous compounding limit, $P_0 e^{rt}$.

Answer to Exercise 23 (on page 53)

Each year, $\frac{7}{8}$ of the employees stay, so after t years there are $480,000(0.875)^t$. The spreadsheet starts 480,000, 420,000, 367,500, 321,562, ... and ends with about 96,680 after 12 years.

Half the company is left when $(0.875)^t = 0.5$:

$$t = \frac{\log 0.5}{\log 0.875} \approx 5.19$$

So the company first drops below half its size during year 6 (after 6 years, 215,422 remain).



INDEX

- arithmetic sequence, 32
- atmospheric pressure, 13
- barometer, 19
- barometric formula, 15
- barometric pressure, 13
- categorical, 4
- categorical variable, 4
- comma-separated values, 5
- common difference, 32
- common ratio, 35
- complex conjugate, 27
- complex number, 25
- Complex Numbers, 25
- complex numbers
 - addition, 26
 - multiplication, 26
- complex plane, 26
- compound interest, 45, 50
- compound interest formula, 47
- conjugate, 27
- continuous compound interest, 49
- continuous compounding, 49
- DataFrame, 4
- datatypes, 3
- e, 44
- explicit formula, 31
- exponent, 39
- exponential decay, 52
- exponential function, 45
- exponential growth, 47, 50
- exponents, 39
 - base, 39
 - fractions, 41
 - negative, 40
 - product rule, 40
 - zero, 40
- Fibonacci sequence, 34
- fundamental theorem of algebra, 28
- future value, 59
- geometric sequence, 35
- half-life, 52
- imaginary number, 25
- imaginary part, 25
- imaginary unit, 25
- index, 5
- ln, 44
- log, 41
 - in python, 42
- logarithm, 41
 - change of base, 44
 - identities, 43
 - in spreadsheets, 44
 - natural, 44
 - product rule, 43
 - quotient rule, 43

magnitude, [26](#)
millimeters mercury, [19](#)

nth root, [41](#)
null, [3](#)

P-trap, [22](#)
pandas, [4](#)
 index, [5](#)

pascal, [13](#)
pascals, [13](#)
pressure, [13](#)
primary key, [3](#)
principal, [47](#)

radioactive decay, [51](#)
real part, [25](#)
recursive formula, [31](#)

sequence, [31](#)
Series, [4](#)
siphon, [21](#)
straw
 drinking, [16](#)