Chapter 2

Frequency and pitch

We saw in Chapter 1 that sound consists of a wave of pressure moving through air at a speed of 343 m/s. The characteristics of the sound we hear are determined by the waveform, the way the sound pressure varies over time. Almost any sound can be used to create music, but most music focuses on a particular subset of sounds, those with periodic waveforms, which correspond to musical notes with clear pitch. In this chapter we look in detail at periodic waveforms and the elementary building blocks of tonal music: notes, pitch, and scales.

2.1 Periodic waveforms and frequency

Listen to the sound file `clicks.mp3`. It contains a recording of a series of clicks, one after another. The clicks start slowly and initially they are easily distinguished from one another. Over time, however, they get faster and faster until at some point they blend together into a continuous sound. Instead of hearing separate individual clicks, you hear a continuous tone. As the clicks come faster still, the pitch of the tone rises, starting low and ending up very high. The shape of the waveform for this particular sound is shown in Fig. 2.1. Each click corresponds to a spike in the sound pressure and the spikes start far apart but get closer together over time.

File: `clicks.mp3`

This experiment reveals a general truth about human perception: if individual events happen slowly then we perceive them as separate, but if they happen sufficiently quickly then they blend together into a continuous sensation. This is true of sound, as here, but it's also true for instance of vision. Repeated flashes of light, as from a stroboscope, will be perceived as separate if they happen slowly, but will blend together into a continuous illumination if they go fast enough. This phenomenon is central, for example, to the way film and video work. Video consists of many individual frames of photography shown in rapid succession. If the frames come slowly we see them as separate and you get a jerky "slow motion" effect. But if they are fast enough they blend together and we see a smoothly moving picture. The point

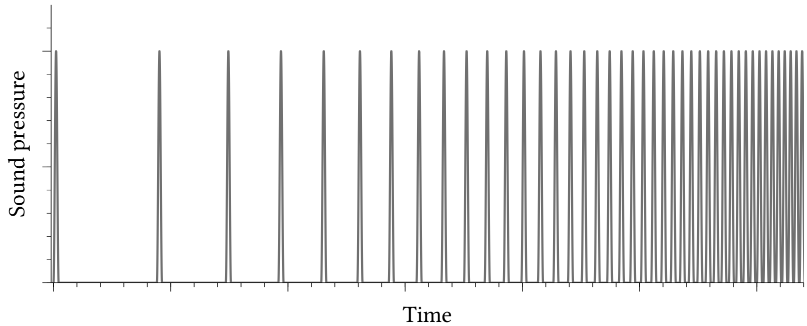

Figure 2.1: A sketch of the waveform for the accelerating series of clicks heard in the sound file `clicks.mp3`.

at which frames or flashes of light start to blend together is called the *flicker fusion rate*. It varies somewhat from one person to another, but for most people it is between about 20 and 40 frames per second. This is precisely why TV, for example, operates at 30 frames per second. If it were much slower than that, people would see separate frames instead of a moving picture. Interestingly, the flicker fusion rate varies not only among people but among different species as well, as shown in Fig. 2.2. It can be as slow as 15 frames per second for sea turtles, as fast 80 per second for dogs, or an amazing 250 per second for fruit flies. TV shows probably don't look very convincing to dogs: their vision is fast enough that they would see the jerky motion of separate frames rather than the smooth moving pictures that humans see.

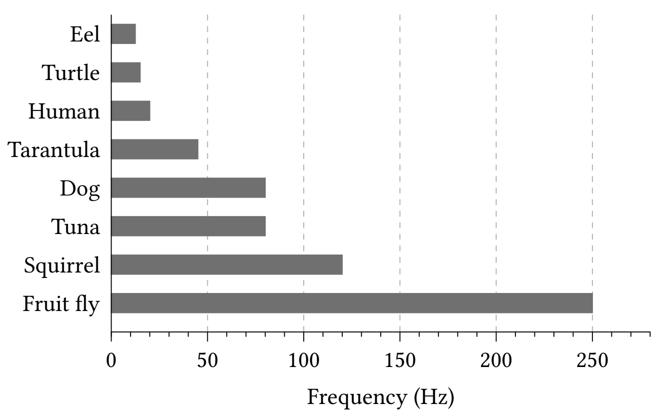

Figure 2.2: Flicker fusion rate for various species. Data from Healy *et al.* (2013).

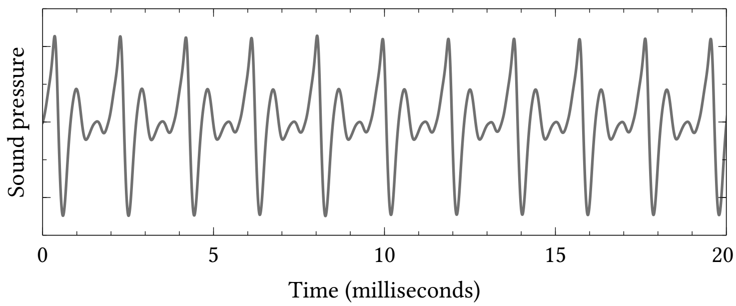

Figure 2.3: Trumpet waveform. The measured sound waveform of a single note played on a trumpet. Note the periodicity of the waveform—it repeats the same shape over and over. The rate of the repeats—the frequency of the waveform—determines the pitch of the note we hear.

A similar thing happens with the clicks in our sound file. If they are faster than about 20 per second we hear a continuous sound and not separate clicks. It is interesting that at 20 clicks per second the fusion rate for sound is about the same as it is for vision, although it's not clear whether this is some fundamental principle of human neuroscience or just a coincidence. At any rate, waveforms that repeat slower than about 20 times a second we hear as individual sounds; those that repeat faster we hear as a continuous tone.

In some cases we may not hear the slower waveforms at all. As we will see in Section 3.6.1, the human ear is quite insensitive to low-frequency sounds.

The sound wave in this experiment is an example of a periodic waveform, one that repeats the same shape over and over again. The shape in this case is just a simple spike but periodic waveforms can be more complicated than this. Figure 2.3 shows the waveform of a note played on a trumpet. As we can see, though the shape of the waveform is quite complex, it has a clear periodicity: the sound pressure repeats the same pattern of variation over and over again.

Not all sound waves are periodic. There are many that are not. But periodic ones are particularly important in music because they correspond to continuous tones that have a distinct pitch—what we call "notes." Non-periodic waveforms, on the other hand, have no clear pitch and do not sound like notes according to the accepted definition of the word.

Familiar examples of non-periodic sounds include cymbal crashes, the sound of running water, and the hissing sound we call "white noise."

A crucial feature of a periodic waveform is its *frequency*. The frequency is the number of times the waveform repeats per second. For instance, looking at Fig. 2.3 we see that the waveform repeats about 10 times over the course of 20 milliseconds or 0.02 seconds. To calculate the frequency (traditionally represented by the letter f), we divide the number of repeats by the amount of time they take in seconds. So in this case

$$f = \frac{10 \text{ cycles}}{0.02 \text{ seconds}} = 500 \text{ cycles per second.} \tag{2.1}$$

This is just a rough calculation. A more careful measurement of the waveform in the figure shows that the real frequency is slightly higher than this, at 523 cycles per second. In scientific parlance, cycles per second are also called *hertz* (named after the 19th century German physicist Heinrich Hertz), so we can also say that the frequency is 523 hertz, or 523 Hz for short.

An alternative way to describe the rate at which the waveform repeats is to specify its *period*—the amount of time taken for one complete cycle of the wave. The period, often represented by the letter T, is equal to $1/f$. For instance, if the frequency is 100 cycles per second, then one cycle takes a hundredth of a second or $1/100$ seconds. So in general the period is given by the formula

$$T = \frac{1}{f}\,. \tag{2.2}$$

2.1.1 Wavelength

Another way to describe a periodic waveform is its *wavelength*. Figure 2.4 shows a sketch of a person listening as a sound wave travels past them. The wavelength of the sound is the distance in space spanned by one complete cycle of the wave as it travels along, usually measured in meters and denoted by the Greek letter λ ("lambda"). If we know the frequency of the sound then it is straightforward to calculate its wavelength, because a 1-wavelength portion of the wave, as indicated in the figure, produces one cycle of the sound as it passes the listener's position, marked by the dashed line. We know that one cycle takes an amount of time T, the period of the wave, and we also know that the wave is traveling at the speed of sound $c = 343\,\text{m/s}$, which means that the length of wave that passes the listener in time T is cT. Thus the wavelength is

$$\lambda = cT. \tag{2.3}$$

Equation (2.2) tells us that $T = 1/f$, so this result can also be written as

$$\lambda = \frac{c}{f}. \tag{2.4}$$

You may also sometimes see this rearranged into the forms

$$\lambda f = c \tag{2.5}$$

or

$$f = \frac{c}{\lambda}, \tag{2.6}$$

the latter being useful when we know the wavelength and want to calculate the frequency.

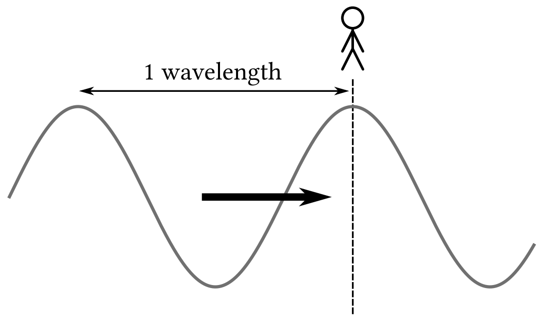

Figure 2.4: Wavelength of a sound wave. The wavelength λ of a periodic sound wave is the length of the space occupied by one cycle of the wave as it travels through the air. By definition this length travels past the listener in one period T of the wave. Since the wave travels at the speed of sound c this means that $\lambda = cT$, or equivalently $\lambda = c/f$.

As an example, consider again the trumpet waveform shown in Fig. 2.3, which has frequency $f = 523\,\text{Hz}$. The wavelength of this sound is

$$\lambda = \frac{c}{f} = \frac{343}{523} = 0.656 \text{ meters}, \tag{2.7}$$

or 65.6 cm, which is a typical wavelength for a musical waveform.

One thing to notice about this number is that it is comparable with the sizes of many everyday objects. Tables, chairs, doorways, windows, musical instruments, and many other things are around this size, give or take a bit. This has important implications for the way sound interacts with those objects, how it bounces around rooms, or is blocked by obstructions. We will see why this is in Section 3.3.

2.1.2 Pitch

When a periodic sound waveform hits our ears, we hear a musical note. The pitch of the note is determined by the frequency of the waveform. Higher frequencies correspond to higher notes and lower frequencies to lower notes. Let us see how this correspondence works.

Figure 2.5 shows a simple experiment to determine the frequency of a note, which can be done without any complicated equipment. A large cog wheel with teeth around its edge is mounted vertically. It is better if the wheel is heavy—made of metal, say—to make it easier to turn steadily. We take a piece of card or some similar material and rest it against the teeth at the top of the wheel, then we start turning the wheel, for instance using a handle or an electric motor. Every time the card hits one of the teeth it makes a small click and the overall resulting sound has a periodic waveform: if the wheel is turning reasonably fast you'll hear a buzzing "note" coming from the piece of card.

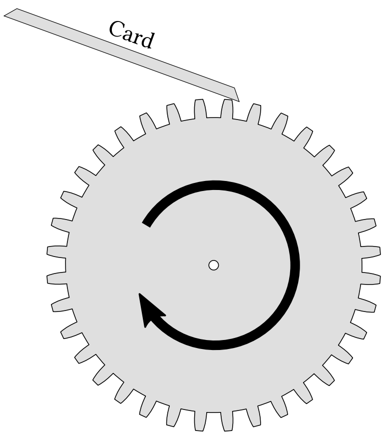

Figure 2.5: Experiment to measure frequency. In this experiment a toothed wheel is turned at a steady pace so that the piece of card at the top makes a buzzing note as it hits the teeth. If we know the rate of rotation of the wheel and the number of teeth we can calculate the frequency of the note.

Now we do the following. First, we play a particular note on a musical instrument. Say we play the note "middle C" on a piano. (If you're not familiar with musical pitches we'll see shortly what middle C means, but for now all you need to know is that it is one of the notes in the middle of the piano keyboard, not too high and not too low.) Now we spin up the wheel until the note it makes matches the one we hear on the piano, so the wheel is also making the note middle C, or as close as we can get.

We can now calculate the frequency of this note as follows. First we count how many times the wheel turns over the course of a minute. Suppose it turns 150 times. Then we stop the wheel and count how many teeth it has. Suppose it has 100 teeth. This means that during the course of a minute the number of times a tooth hits the piece of card is $150 \times 100 = 15\,000$. This is the number of cycles of the sound waveform in a minute. To get the number of cycles in a second we divide this figure by 60, so the frequency of the sound is

$$f = \frac{15\,000 \text{ cycles}}{60 \text{ seconds}} = 250 \text{ cycles per second}, \tag{2.8}$$

or 250 Hz.

This is actually not far off the correct result, although it's not perfect: an accurate measurement of the frequency of middle C would find that the correct number is 261.6 Hz. If you actually do this experiment, or something like it, you will probably find that you also don't get exactly the correct result, in part because it's difficult to hold the speed of the wheel steady. But more sophisticated versions of the same experiment, for instance using an electrically operated wheel, or using an electronic signal generator in place of the wheel, can be used to make very accurate measurements of the frequency of any note.

Now suppose we try the experiment again with a different note. Instead of middle C, we use the note "high C," which is an octave higher on the piano keyboard. (Again, if you're not familiar with musical notes or don't know what an octave is, not to worry—we'll cover that shortly.) Repeating the experiment, we now find that we have to turn the wheel twice as fast to match the sound of the higher note and hence it turns about 300 times in a minute instead of only 150 times. Given again that there are 100 teeth, this means there are 30 000 cycles of the sound in a minute, and the frequency is

$$f = \frac{30\,000 \text{ cycles}}{60 \text{ seconds}} = 500 \text{ Hz}. \tag{2.9}$$

In other words, by going up an octave in pitch we have doubled the frequency. This is not a coincidence; it is a general rule. If the pitch goes up by an octave the frequency doubles. This is our first example of a general principle of frequency and pitch: changes in musical pitch correspond to *multiplying* the frequency by some number (in this case 2).

Experiments like this long ago established accurate values for the frequencies of all musical notes. Among other things, we find that the musically useful part of the sound frequency range extends from about 20 Hz to about 4000 Hz. The lowest note you can play on a tuba, for example, clocks in at a frequency of 21.8 Hz, while the highest note on a piccolo has frequency 4186 Hz. The human ear can hear frequencies higher than this, up to almost 20 000 Hz, but those frequencies are not normally used as musical notes and can indeed be quite unpleasant to hear.

2.2 Intervals and scales

If we play two different notes, say on a piano, either at the same time or one after another, the gap between their pitches is called an *interval*. Two notes close together would have a small interval between them; two notes far apart would have a large one.

The first important scientific study of musical intervals is often attributed to Pythagoras, the ancient Greek philosopher-scientist famous for the theorem about triangles. It is unclear whether the study was really performed by Pythagoras or by one of his followers, but it certainly took place and the results are central to our understanding of how music works.

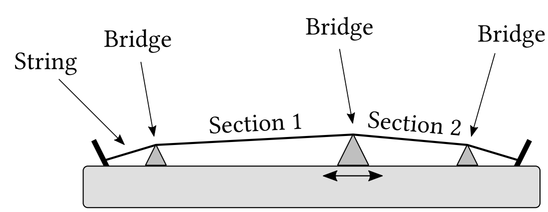

Figure 2.6: A simple one-stringed musical instrument that plays two notes. The sections of string labeled "Section 1" and "Section 2" vibrate when plucked and produce two different notes. The center bridge is movable, and if it is positioned so that the length of one section is a whole-number multiple of the other then it is found that the notes sound pleasing together.

Figure 2.6 shows a version of the experiment Pythagoras is rumored to have performed. It makes use of a simple stringed instrument, sometimes called a *monochord*, which has just one string, similar perhaps to a one-stringed guitar or zither. The string is stretched between two fixed edges, called bridges, with a third, movable bridge in between. You can pluck the two halves of the string, to the left and right of the center bridge, and the instrument produces notes. Typically the two halves produce different notes with different pitches, and moreover both pitches change when the center bridge is moved from side to side. It is found that a shorter segment of string produces a higher note and a longer segment produces a lower one.

We will look at the workings of string instruments in Chapter 10.

Now one moves the position of the central bridge back and forth, trying out different positions, plucking the two halves of the string at the same time in each position and listening to the musical interval between the two notes produced. Some intervals sound pleasant or *consonant* to the human ear. The notes go well together. Other intervals sound unpleasant or *dissonant*. The crucial discovery that Pythagoras made (if indeed it was him) was that when the length of one half of the string is a simple whole-number multiple of the length of the other—like twice as long or three times as long—then the interval they make sounds pleasant to our ears.

We study the science behind dissonant and consonant intervals in more detail in Section 5.4.3.

With the benefit of modern scientific insight, we can today understand what this means: if the lengths of the sections of the string have a certain ratio, then the frequencies of the notes they make have the same ratio. For instance, if the lengths have a ratio of 2-to-1 then one note has twice the frequency of the other. So what Pythagoras in effect discovered was that if the frequency of one note is a simple multiple of the other then the two notes sound pleasant together. This fundamental discovery is central to the creation of harmonious sounding music and forms the foundation for scales and harmonies in virtually every musical tradition worldwide.

The connection between the length and frequency of a string is governed by Mersenne's law, which is discussed in Section 9.2.1.

2.2.1 Frequencies and intervals

The simplest example of the principle described in the previous section occurs when the frequency of one note is twice that of the other. When this happens, the two notes produce the musical interval known as an *octave*. Indeed this is the very definition of an octave: it is the sound made by two notes when one has twice the frequency of the other. The best way to understand what an octave sounds like is to hear it. Listen to the sound file `octave.mp3` for an example. If you can't play the sound file, think instead of the first two notes of the tune *Over the Rainbow* from the movie *The Wizard of Oz*, which starts with the words "Somewhere over the rainbow". The two notes sung to the first word "somewhere" form an octave.

File: `octave.mp3`

The octave is a little difficult to describe in words, but perhaps the nearest one can come to it is to say that when you start with a note and go up an octave you get another note that sounds "the same but higher." It has the same musical quality as the first note, just higher in pitch. This phenomenon of sameness between notes an octave apart is known as *octave equivalence*.

The octave is universally acknowledged to be the most consonant of musical intervals. It is a fundamental element of harmony in virtually every musical tradition and pretty much everyone agrees that two notes an octave apart sound pleasing together. The octave also provides us with one of the simplest examples of a general law of music that we touched upon earlier: when the pitch goes up by a certain interval you *multiply* the frequency by a number. You don't add to the frequency, you multiply it—by two in the case of the octave. Conversely, when the pitch goes down you divide the frequency by a number. For instance, to go down by an octave you divide by two.

As Pythagoras discovered, multiplying by other whole numbers gives us other pleasing intervals. Multiplying by three gives us a larger pitch jump of more than an octave that, for reasons we will see shortly, is called a twelfth, or perhaps more commonly "an octave plus a fifth." (Listen to the sound file `twelfth.mp3` to hear an example.) The intervals of an octave and an octave plus a fifth are depicted visually in Fig. 2.7. An octave plus a fifth is a large interval. We can reduce it to a more manageable size by taking the note down an octave, which gives you the musical interval known as a fifth (or more properly a "perfect fifth"). Thus, when we multiply by three we go up an octave plus a fifth and when we go down an octave again—which means dividing by two—we get a fifth. To put that another way, to go up a fifth you multiply by three and divide by two, or equivalently you multiply by $\frac{3}{2}$.

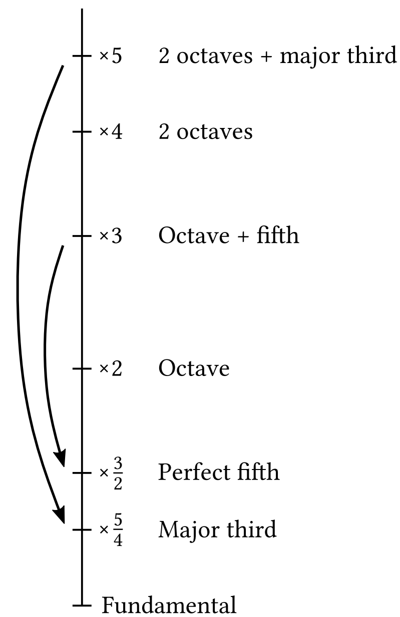

Figure 2.7: A visual depiction of the frequencies of notes in a major chord. Starting at any note ("the fundamental"), multiplying the frequency by 2, 3, 4, or 5 gives us notes that are, respectively, an octave higher, an octave plus a fifth, two octaves, and two octaves plus a major third. Shifting down an octave from an octave plus a fifth gives us the "perfect fifth," which has frequency $\frac{3}{2}$ times the fundamental. Similarly the major third has frequency $\frac{5}{4}$ times the fundamental.

The fifth is perhaps the second most universal and important interval in music, also widely agreed to have a pleasing sound. Going up a fifth corresponds to multiplying the frequency by $\frac{3}{2}$ and going down a fifth corresponds to dividing by $\frac{3}{2}$. You can hear the sound of a fifth in the sound file `fifth.mp3`. A simple example of a fifth in music is the start of the tune *Twinkle Twinkle Little Star*—the second "twinkle" is a fifth higher than the first one.

We can continue this process. If you multiply the frequency by four, the pitch goes up by an interval of two octaves (and if you divide by four it goes down two octaves). This is logical: we have said that to go up by one octave you multiply by two, so to go up two octaves you multiply by two then two again, which is the same as multiplying by four.

If you multiply the frequency by five, the pitch goes up by more than two octaves—by two octaves and a major third in fact. Again this is a large jump in pitch but we can bring it down to a more reasonable range by going down two octaves, which gives us the musical interval known as a major third. Going down two octaves involves dividing by four, so in other words the interval of a major third corresponds to multiplying by five and dividing by four, or equivalently multiplying by $\frac{5}{4}$.

Files:
`twelfth.mp3`
`fifth.mp3`
`majorthird.mp3`

Now we have four notes that all sound pleasing together: the first note we started with, the third, the fifth, and the octave. These notes together form a major chord, the most fundamental building block of Western musical harmony. You can hear an example in the sound file `majorchord.mp3`.

File: `majorchord.mp3`

2.2.2 The major scale

The basic principle at work here is that pairs of notes sound pleasing together if the frequency of one is a simple multiple of the other. Simple multiples in this case include basic whole numbers—two times, three times, four times—but also small fractions like $\frac{3}{2}$ or $\frac{5}{4}$. By applying these ideas we can quickly find other notes that also sound good together.

For example, we have seen that the octave and the fifth are two of the most fundamental intervals. If we combine them by going up an octave and then down a fifth, we get the musical interval known as a (perfect) fourth. Going up an octave means multiplying the frequency by two and going down a fifth means dividing by $\frac{3}{2}$, so a fourth is equivalent to multiplying by

$$\frac{2}{3/2} = 2 \times \tfrac{2}{3} = \tfrac{4}{3}. \tag{2.10}$$

Similarly, going up a fifth and down a fourth gives the musical interval of a major second, with a multiplier of

$$\frac{3/2}{4/3} = \tfrac{3}{2} \times \tfrac{3}{4} = \tfrac{9}{8}. \tag{2.11}$$

Going up a third, then up again either a fourth or a fifth, gives us the intervals known as the major sixth and major seventh, with multipliers

$$\tfrac{5}{4} \times \tfrac{4}{3} = \tfrac{5}{3}, \tag{2.12}$$
$$\tfrac{5}{4} \times \tfrac{3}{2} = \tfrac{15}{8}. \tag{2.13}$$

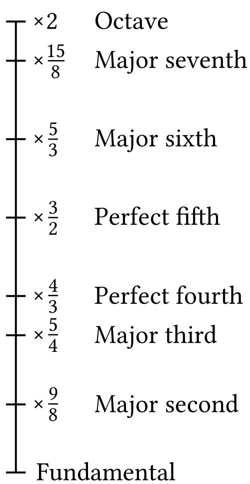

Figure 2.8: Frequencies of the notes in the major scale. Starting with any fundamental frequency for the first note of the scale, multiplying by each of the fractions here produces the notes of the major scale.

Put all of these different notes together and we get the set of eight notes depicted in Fig. 2.8. These eight notes form the *major scale*, one of the fundamental building blocks of Western music. The major scale consists of the starting note, called the *fundamental* or *tonic*, and seven other notes whose frequencies are all simple multiples of the fundamental. Because they are simple multiples the notes all sound good together. They provide a guaranteed pleasing sound that makes it easy to create harmonious music: essentially any combination of the notes of the major scale will sound pleasant, which is precisely why they play such a large role. You can hear a major scale by playing the white keys on a piano (which play the C major scale, i.e., the scale whose fundamental is the note C) or by listening to the sound file `majorscale.mp3`.

File: `majorscale.mp3`

2.2.3 Problems with the major scale

The major scale of Section 2.2.2 produces a pleasing sound and is based on solid mathematical principles. You might be surprised then to learn that in fact this scale is hardly ever used. The "major scale" as we know it is not this set of notes, but a slightly different one.

To understand why this is, consider playing the scale starting on the note middle C—playing in the *key of C major*, as a musician would say. As we saw in Section 2.1.2, middle C has a frequency of 261.6 Hz. To form a major scale starting on middle C we would multiply this fundamental frequency by each of the fractions in Fig. 2.8 in turn. For instance, to get the frequency of the second note of the scale, which is D, we would multiply by $\frac{9}{8}$ thus:

$$f_{\mathrm{D}} = \tfrac{9}{8} \times 261.6 = 294.3\,\mathrm{Hz}, \tag{2.14}$$

and to get the sixth note of the scale, which is A, we would multiply by $\frac{5}{3}$ thus:

$$f_{\mathrm{A}} = \tfrac{5}{3} \times 261.6 = 436.0\,\mathrm{Hz}. \tag{2.15}$$

But now suppose we play in the key of D major, meaning that we start our major scale on the note D, which has frequency 294.3 Hz as shown above. Then we would multiply this latter number by the fractions in Fig. 2.8 to get the notes of the scale. For instance, to get the fifth note of the scale, which is A, we would multiply by $\frac{3}{2}$ thus:

$$f_{\mathrm{A}} = \tfrac{3}{2} \times 294.3 = 441.5\,\mathrm{Hz}. \tag{2.16}$$

But now we notice a disturbing fact: the frequency of the note A calculated as part of the C major scale in Eq. (2.15) is not the same as the frequency of the note A as part of the D major scale in Eq. (2.16).

There is nothing wrong with our calculations. The numbers are correct. This is a real feature of the scale: the frequency of a particular note depends on what key you are playing in. If you are playing in C major the note A has one frequency. If you are playing in D major it has another.

The difference between the two frequencies is known in musical theory as a *comma*. This particular comma is technically called the *syntonic comma.*

This is a serious problem. It means for instance that if you wanted to play the piano in the key of D you would have to tune it differently from the way you tune it in the key of C. You would have to retune the entire piano—a laborious task taking hours—every time you wanted to play in a new key. This of course is entirely impractical and would make changing from one key to another on an instrument like the piano essentially impossible. For many centuries composers and musicians got around this problem by avoiding it altogether and only playing a very limited range of keys and scales. This is part of the reason why ancient (pre-Renaissance) music has a simpler harmonic structure than modern music. In more recent times, however, a new solution to the problem was discovered, the so-called equal temperament scale, which allows musicians to play in any key without retuning their instruments. Modern Western music owes essentially its entire existence to this crucial innovation, which relies on an amazing mathematical coincidence, as we'll see in the next section.

2.3 Equal temperament tuning

Equal temperament tuning provides an alternative way to define note frequencies, different from the scheme described above in Section 2.2.2 which is called *just tuning* or *just intonation*. Equal temperament produces notes with slightly different pitches from those of just intonation, but the difference is small enough that in most cases one does not notice it.

The equal temperament scheme works by dividing an octave into twelve equally sized pitch steps, which for historical reasons are called *half-steps* by musicians, or sometimes *semitones*. As we have said, going up by a certain interval involves *multiplying* the frequency of a note by some number. What number do we multiply by to go up a half-step? We can work out the answer as follows.

Let us write the number we want as x, an unknown quantity. There are by definition twelve half-steps in an octave, all of equal size, so if we multiply by x twelve times we go up an octave. At the same time we know that going up an octave means multiplying by two. Hence multiplying by x twelve times is the same as multiplying by two. In mathematical terms, $x^{12} = 2$. Rearranging this equation for x, we find that

$$x = 2^{1/12} = 1.05946. \tag{2.17}$$

In the equal temperament system this result defines the frequencies of the notes. You start at any note you like—say middle C—and from there you can go up a half-step by multiplying by $2^{1/12}$. Or you can go up by two half-steps, also called a *whole-step*, by multiplying by $2^{1/12}$ twice, which means multiplying by $2^{1/12} \times 2^{1/12} = 2^{2/12}$. Similarly, going up three half-steps means multiplying by $2^{3/12}$, four half-steps means $2^{4/12}$, and so forth. To go up by n half-steps you multiply n times by $2^{1/12}$, which is

$$\underbrace{2^{1/12} \times 2^{1/12} \times \ldots \times 2^{1/12}}_{n \text{ times}} = 2^{n/12}. \tag{2.18}$$

To put that another way, if note 1 with frequency f_1 is n half-steps higher than note 2 with frequency f_2, then

$$f_1 = 2^{n/12} f_2. \tag{2.19}$$

We can use this equation to find the frequency of any higher note given the frequency of a lower note. This in turn means that if we fix the frequency of any single starting note we fix the frequencies of all higher notes as well. Similarly, to go down by n half-steps we divide by $2^{n/12}$, which means

$$f_1 = \frac{f_2}{2^{n/12}}, \tag{2.20}$$

and hence all lower notes are fixed as well.

Alternatively, suppose we are given a higher frequency f_1 and lower frequency f_2 and we want to know how many half-steps there are between them. Dividing both sides of Eq. (2.19) by f_2 and taking the logarithm, we get

$$\log \frac{f_1}{f_2} = \log 2^{n/12} = \frac{n}{12} \log 2, \tag{2.21}$$

and rearranging for n then gives

$$n = \frac{12}{\log 2} \log \frac{f_1}{f_2}. \tag{2.22}$$

This equation tells us how many half-steps there are between any two frequencies we specify. This equation will come in handy many times in this book.

2.3.1 The equal temperament major scale

In the equal temperament system the major scale is defined completely differently from the way we defined it in Section 2.2.2. The major scale starting on any given note—any fundamental—is defined as the set of pitches you get by going up 2, 4, 5, 7, 9, 11, and 12 half-steps from that fundamental. The amazing thing about this seemingly random set of numbers is that it gives *almost exactly* (but not quite) the just intonation major scale that we derived from first principles before. Here's what the frequencies look like:

		Frequency multiplier	
Note	half-steps	Equal temperament	Just intonation
Fundamental	0	1	1
2nd	2	$2^{2/12} = 1.122$	$\frac{9}{8} = 1.125$
3rd	4	$2^{4/12} = 1.260$	$\frac{5}{4} = 1.250$
4th	5	$2^{5/12} = 1.335$	$\frac{4}{3} = 1.333$
5th	7	$2^{7/12} = 1.498$	$\frac{3}{2} = 1.500$
6th	9	$2^{9/12} = 1.682$	$\frac{5}{3} = 1.667$
7th	11	$2^{11/12} = 1.888$	$\frac{15}{8} = 1.875$
8th	12	2	2

As you can see, the frequency multipliers for each of the eight notes of the major scale, as calculated using equal temperament, are very close to the "perfect" values that we get from just intonation. Recall that the latter are based on the idea that simple multiples of frequencies sound good to our ears. Equal temperament tuning does not produce simple multiples, but it almost does. For instance, the note called the fifth (because it is the fifth note of the scale) should have frequency $\frac{3}{2}$ or 1.5 times the frequency of the fundamental. In the equal temperament scale it is actually 1.498 times the fundamental. There is a difference between the two, but it is very small—so small that the human ear cannot hear it. Similarly the fourth should be 1.333 times the fundamental frequency, but in the equal temperament scale it is 1.335. In essence the equal temperament scale is a very slightly out-of-tune version of the just intonation major scale.

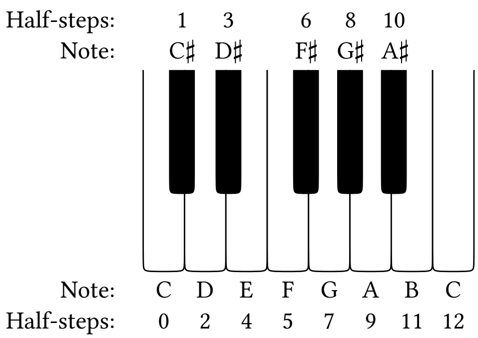

Figure 2.9: One octave of piano keys. The white keys on the piano make up the notes of the C major scale, which start at C and go up by 2, 4, 5, 7, 9, 11, and 12 half-steps as shown. This leaves five other notes at 1, 3, 6, 8, and 10 half-steps from the fundamental, which are played using the black keys and produce the notes C♯, D♯, F♯, G♯, and A♯.

The remarkable thing is that *all* the notes of the scale are well approximated by notes in the equal temperament system. There is no reason why this has to be the case. It just so happens that if you divide an octave into 12 equally sized parts, giving 13 notes including the starting and ending notes, then 8 of those 13 notes *very nearly* make a just-intonation major scale. This is purely a mathematical coincidence, but it is a very useful one—one on which all of modern Western music rests. In practice, modern music rarely uses the just intonation scale because even though just intonation gives perfect intervals that our ears hear as pleasing sounds, it is impractical, requiring us to retune our instruments every time we play in a different key. Instead we use the equal temperament scale, which is slightly out of tune but, as we will see in a moment, doesn't require us to retune any instruments.

2.3.2 The 12-tone scale

The notes of modern Western music are defined using the equal temperament tuning system and referred to by letter names starting from the beginning of the alphabet: A, B, C, and so forth. Two notes an octave apart are given the same letter name. If we go up an octave from an A we reach another, higher A. Thus "A" is the name of many different notes, not just one. This might seem confusing at first, but it has some advantages. As we have said, two notes an octave apart sound similar to human ears, so giving them the same name is somewhat logical.

Consider the major scale again, which has eight notes. As we have said, one way to get a major scale is to play the white keys on a piano. If you start at the right place on the keyboard (and it does have to be the right place) and if you play the eight white keys shown in Fig. 2.9 from left to right, you will hear a major scale. Starting with the lowest note (on the left end) these eight notes have the letter names C, D, E, F, G, A, B, and C again, the last note being an octave up from the first, so it gets the same letter name.

Why do the notes of the scale start at C, you ask? Why not A? Why indeed. It doesn't make much sense. But it's the way things have been done for hundreds of years and it's too late to change it now. So instead of A, B, C, the scale starts C, D, E. There are eight notes in the scale, but the first and last have the same letter name C, so we only need seven letters to label all eight notes. Thus the letter names go up to G, the seventh letter of the alphabet. Once you reach G, you wrap around to the beginning of the alphabet and letter A again. Hence the sequence C, D, E, F, G, A, B, C.

Another way to think about these notes is in the language of half-steps where, as we have seen, the notes of the major scale start at the fundamental and go up 2 half-steps then 4, 5, 7, 9, 11, and 12. These numbers are also noted in Fig. 2.9. But this now raises a question: what about the other half-steps? What if we go up 1, 3, 6, 8, or 10 half-steps? These intervals give us five additional notes that do not belong to the major scale. They fall half way between scale notes. On the piano these extra notes are produced by the black keys, which are located above and between the white keys as shown in Fig. 2.9. The result is the classic piano-style keyboard arrangement with twelve notes in total, or thirteen if you count the two Cs at the top and the bottom.

Half-steps from C	Note name(s)
0	C
1	C♯ or D♭
2	D
3	D♯ or E♭
4	E
5	F
6	F♯ or G♭
7	G
8	G♯ or A♭
9	A
10	A♯ or B♭
11	B
12	C

Table 2.1: The notes of the 12-tone equal temperament scale, starting on C. There are 13 notes listed because we include two Cs, at the bottom and top of the scale, but there are only 12 distinct note names, hence the name "12-tone" scale.

The additional five notes on the black keys are named after the scale notes closest to them. The one a half-step up from C is called C♯, pronounced "C sharp." The one a half-step up from D is D♯ and the others are F♯, G♯, and A♯. Just to make things a little more complicated, however, the black keys also have another set of alternate names. The note a half-step *down* from D is called D♭, pronounced "D flat," which is the same note as C♯. So D♭ and C♯ are two different names for the same thing. Similarly the other four notes are called E♭, G♭, A♭, and B♭. You should be familiar with both sets of names, as both are commonly used by musicians.

The end result is that the octave from C to C is divided up into twelve equally sized half-steps. Together they form the *12-tone equal temperament* scale, sometimes abbreviated as 12-TET. This scale is the foundation on which virtually all modern Western music is based (as well as music in many other traditions). The notes of the 12-tone scale are listed in Table 2.1.

Figure 2.9 represents just one octave of notes, but the same motif can be repeated over other octaves. For instance, we can start at the upper of the two Cs and construct another C major scale working upward from there, adding an additional octave on

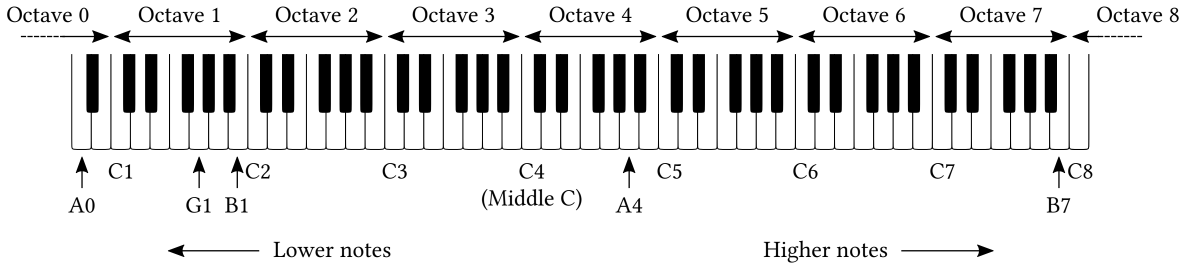

Figure 2.10: The piano keyboard. A full piano keyboard spans seven octaves, numbered from 1 to 7 as shown, plus a few extra notes at the ends that fall in octave 0 and octave 8.

top of what we already have. This next octave is represented by a pattern of 12 keys exactly the same as the one we already have. We can also add lower octaves in the same way, and we can repeat the process as much as we like, adding octave after octave, either above or below. The result is a full piano keyboard as shown in Fig. 2.10.

The standard piano keyboard has seven octaves from top to bottom with the lowest notes on the left end and the highest on the right. Because the notes in each octave have the same letter names it can be difficult to know which note someone is talking about if they only give the letter name. Some notes, like "middle C," have their own names but most do not. To get around this problem, we number the octaves from 1 to 7, starting at the low end and working up. Each octave starts at C and goes up to the next highest B. Then the next octave starts with the C after that—see Fig. 2.10.

A few high-end concert pianos, particularly those made by the Bösendorfer company, go lower, down to F0 or in extreme cases all the way down to C0. These notes are useful for playing certain music originally written for pipe organ, but are not widely used in the piano repertoire. Indeed, the fundamental frequency of the note C0 is below the normal hearing range for most people.

Thus, for example, the lowest C on the piano keyboard is called C1, meaning the C in octave 1. The G above it is G1, the B above that is B1, but the C above *that* is C2. And so forth. Middle C is technically C4, and the highest complete octave ends with the note B7. The traditional piano keyboard extends one half-step higher than this, to C8, and three half-steps lower than C1 at the left end of the keyboard, to the notes B0, B♭0, and A0 (in the "zeroth" octave). The range of notes can be extended still further in principle, up into octave 9 or 10, or down into octave −1 or −2, though in practice this is done rarely, if ever.

2.4 Frequencies of notes

Here is an important point: if we specify the frequency of just one single note on the piano keyboard, then we can immediately calculate the frequency of every other note unambiguously. In the equal temperament system, every note has a single unique frequency that can be calculated using the formulas given in Section 2.3. This contrasts with the just intonation scale of Section 2.2.3, for which the frequency of the same note can be different when playing in different keys. In the equal temperament scale there is just one frequency for every note. This is the primary reason why we use equal temperament. Equal temperament is in some ways a compromise, with notes that are slightly out of tune, but the convenience of having one frequency for every note outweighs other considerations.

The calculation of note frequencies is straightforward. We start off by fixing the frequency of one note. By convention that note is A4, the A in the middle of the piano keyboard (see Fig. 2.10), also sometimes called "tuning A." By international standard, established in 1955, the frequency of A4 is 440 Hz exactly. We write

$$f_{\mathrm{A4}} = 440\,\mathrm{Hz}. \tag{2.23}$$

There is no scientific principle behind this. It is just a number picked by a committee for convenience. But it is widely used the world over as the standard by which musical notes are defined.

Once we have fixed the frequency of this one note, we can calculate the frequency of any other using Eqs. (2.19) and (2.20). For instance, when we go up one half-step we multiply the frequency by $2^{1/12} = 1.05946$. Going up a half-step from A4 takes us to B♭4, so the frequency of B♭4 is

$$f_{\mathrm{B\flat 4}} = 440 \times 2^{1/12} = 466.2\,\mathrm{Hz}. \tag{2.24}$$

Similarly, going up two half-steps from A4 to B4 gives us a frequency of

$$f_{\mathrm{B4}} = 440 \times 2^{2/12} = 493.9\,\mathrm{Hz}, \tag{2.25}$$

and going up three half-step to C5 gives us

$$f_{\mathrm{C5}} = 440 \times 2^{3/12} = 523.3\,\mathrm{Hz}, \tag{2.26}$$

and so forth. The general formula for the note n half-steps up from A4 is

$$f = 440 \times 2^{n/12}. \tag{2.27}$$

Given this formula we can calculate the frequencies of the notes all the way to the top of the piano keyboard, or beyond.

Similarly, to go down a half-step from A4 we divide by $2^{1/12}$, giving the frequency of A♭4 to be

$$f_{\mathrm{A\flat 4}} = \frac{440}{2^{1/12}} = 415.3\,\mathrm{Hz}. \tag{2.28}$$

Going down two half-steps gives us G4:

$$f_{\mathrm{G4}} = \frac{440}{2^{2/12}} = 392.0\,\mathrm{Hz}, \tag{2.29}$$

Note	Frequency (Hz)
A0	27.50
A#0	29.14
B0	30.87
C1	32.70
C#1	34.65
D1	36.71
D#1	38.89
E1	41.20
F1	43.65
F#1	46.25
G1	49.00
G#1	51.91
A1	55.00
A#1	58.27
B1	61.74
C2	65.41
C#2	69.30
D2	73.42
D#2	77.78
E2	82.41
F2	87.31
F#2	92.50
G2	98.00
G#2	103.8
A2	110.0
A#2	116.5
B2	123.5
C3	130.8
C#3	138.6
D3	146.8
D#3	155.6
E3	164.8
F3	174.6
F#3	185.0
G3	196.0
G#3	207.7
A3	220.0
A#3	233.1
B3	246.9
C4	261.6
C#4	277.2
D4	293.7
D#4	311.1
E4	329.6
F4	349.2
F#4	370.0
G4	392.0
G#4	415.3
A4	440.0
A#4	466.2
B4	493.9
C5	523.3
C#5	554.4
D5	587.3
D#5	622.3
E5	659.3
F5	698.5
F#5	740.0
G5	784.0
G#5	830.6
A5	880.0
A#5	932.3
B5	987.8
C6	1047
C#6	1109
D6	1175
D#6	1245
E6	1319
F6	1397
F#6	1480
G6	1568
G#6	1661
A6	1760
A#6	1865
B6	1976
C7	2093
C#7	2218
D7	2349
D#7	2489
E7	2637
F7	2794
F#7	2960
G7	3136
G#7	3322
A7	3520
A#7	3729
B7	3951
C8	4186

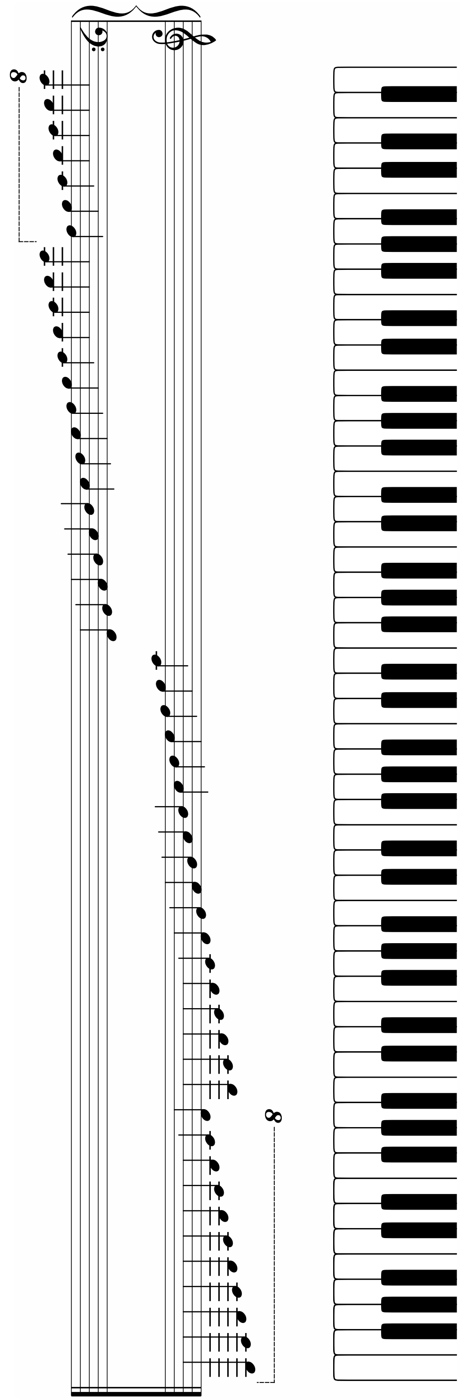

Figure 2.11: Names and frequencies of notes on the piano keyboard. A standard piano keyboard has 88 notes, from a low of A0, with frequency 27.5 Hz, close to the bottom of the hearing range, to a high of C8 at 4186 Hz. A few instruments can go outside of this range—the tuba can go down to E0, for instance—but the seven octaves of the piano represent the basic range of normal musical notes. Below the keyboard we also show the same notes in musical notation. If you are unfamiliar with musical notation, an introduction is given in Section 2.5.

and so forth. The general formula for going down n half-steps from A4 is

$$f = \frac{440}{2^{n/12}}, \tag{2.30}$$

which allows us to calculate the frequencies of notes all the way down to the bottom of the keyboard or beyond. Figure 2.11 lists the frequencies of all the notes on the piano keyboard calculated in this way, and will be a useful reference throughout this book on the many occasions when we will need to know note frequencies.

Because we multiply the frequency by a certain number every time we go up a half-step—as opposed to adding to the frequency, for instance—the frequencies of notes get further apart as we go up the scale. For example, referring to Fig. 2.11, the change in frequency from C5 to C♯5 is 31.1 Hz, but the change from C♯5 to D5 is 32.9 Hz and from D5 to D♯5 is 35.0 Hz, and so forth. Figure 2.12 shows a plot of the frequencies of all the notes on the piano keyboard—the numbers listed in Fig. 2.11—and the growth in the spacing between notes is clearly visible as frequency gets higher. This change in spacing has some practical consequences for music making. For instance, it causes audible problems with the tuning of valve brass instruments such as the trumpet and tuba, requiring modifications to the design of these instruments as well as adjustments in playing style. We discuss these issues in detail in Section 11.4.6.

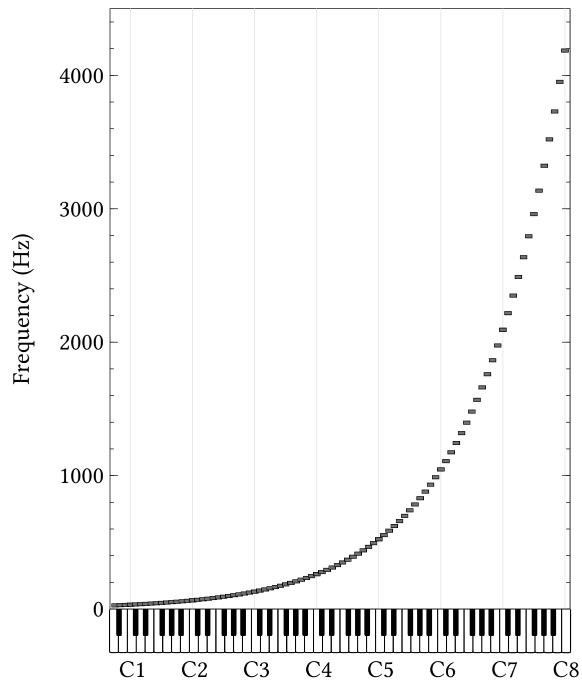

Figure 2.12: The frequencies of all notes on the piano. Each point indicates the frequency of the corresponding note on the piano keyboard at the bottom. Note how the gap between the frequencies of successive notes gets bigger as the notes get higher.

2.4.1 Musical cents

Just intonation and equal temperament differ by only a small amount. The pitches of notes are slightly different in the two schemes, but not by much. In music the smallest difference in pitch that we usually talk about is the half-step or semitone, but if we want to talk about smaller differences we can use another unit: *musical cents.* A musical cent is one hundredth of a half-step and thus there are 100 cents in one half-step. Since there are 12 half-steps in an octave that means there are 1200 musical cents in an octave. We can use this observation to work out the mathematics of musical cents. The calculations are similar to the ones we did for half-steps in Section 2.3.

As we have said, when you increase the pitch of a note the rule is that you *multiply* the frequency by a number, so when the pitch goes up by one musical cent we must multiply by some number. Let us call this number y. We know that there are 1200 cents in an octave, so if we multiply by y 1200 times the pitch will go up an octave. Multiplying by y 1200 times is the same as multiplying by y^{1200}, and going up an octave is the same as multiplying by 2, so $y^{1200} = 2$. Rearranging, we then find that

$$y = 2^{1/1200} = 1.000578. \tag{2.31}$$

In other words, to raise the pitch by one musical cent we multiply the frequency by $2^{1/1200}$. If we want to raise the pitch by n cents then we need to multiply n times by $2^{1/1200}$, which looks like this:

$$\underbrace{2^{1/1200} \times 2^{1/1200} \times \ldots \times 2^{1/1200}}_{n \text{ times}} = 2^{n/1200}. \tag{2.32}$$

A calculation we will often need to do is determine how many musical cents there are between two pitches. Suppose two notes have frequencies f_1 and f_2. If f_1 is n musical cents higher than f_2 then Eq. (2.32) tells us that

$$f_1 = 2^{n/1200} f_2. \tag{2.33}$$

Dividing both sides of this equation by f_2 and taking the logarithm gives

$$\log \frac{f_1}{f_2} = \log 2^{n/1200} = \frac{n}{1200} \log 2, \tag{2.34}$$

and rearranging for n then tells us that

$$n = \frac{1200}{\log 2} \log \frac{f_1}{f_2}. \tag{2.35}$$

This is the number of musical cents between our two frequencies. This formula will be useful to us in many places throughout this book.

One musical cent is a very small difference in pitch, so small that it cannot be heard by the human ear. Listen to the two tones in the sound file `onecent.mp3`, which are exactly one musical cent apart. You should not be able to tell which of the two is the higher note—the difference in pitch is simply too small for even the most sensitive ear to pick up. The smallest pitch difference that anyone can hear is about five cents. Not everyone can hear a pitch difference this small, but those with particularly sensitive ears, including especially musicians who have a lot of experience listening for small differences in tuning, can just make out a pitch difference of

File: `onecent.mp3`

five cents. Listen to the sound file `fivecents.mp3` and see if you can hear the pitch difference between the two notes.

File: `fivecents.mp3`

Larger pitch differences are easier to hear. Almost everyone can hear a pitch difference of 25 musical cents, which is a quarter of a half-step or semitone. A note that is out of tune by 25 cents will sound bad to just about anyone. Listen to the sound file `twentyfivecents.mp3` for an example.

File: `twentyfivecents.mp3`

2.4.2 Comparison between just intonation and equal temperament

How large is the difference between the just intonation major scale and the equal temperament major scale? Is it large enough for the human ear to hear? We can answer this question by working out how many musical cents there are between the notes of one scale and the other. If there is a more than a five-cent difference, then at least some listeners will be able to hear it.

Suppose we start a major scale on a note with frequency f. This is the fundamental or tonic, the first note of the scale. As shown in Fig. 2.8 on page 26, in the just intonation system the next note of the major scale is given by multiplying by $\frac{9}{8}$, so it would have frequency $f_1 = \frac{9}{8}f$. In the equal temperament major scale, on the other hand, the next note is two half-steps up from the fundamental, so it would have frequency $f_2 = 2^{2/12}f$.

These are both versions of the same note, but they are slightly different. How different? We can calculate the number of musical cents between them by plugging the values of f_1 and f_2 into Eq. (2.35), which gives

$$n = \frac{1200}{\log 2} \log \frac{\frac{9}{8}f}{2^{2/12}f} = \frac{1200}{\log 2} \log \frac{\frac{9}{8}}{2^{2/12}} = 3.91 \text{ cents.} \tag{2.36}$$

Notice how the value of the fundamental frequency f has cancelled out of the equation so the final answer does not depend on it. This means that the number of musical cents is the same no matter what note we start the scale on. There will always be 3.91 musical cents between the just intonation and equal temperament versions.

Thus there is indeed a pitch difference between these two notes, but it is very small. 3.91 musical cents is below the five-cent threshold of hearing for even the most sensitive of human ears. No one would be able to hear this tiny pitch difference. This is a good thing. It means that even though equal temperament tuning is a compromise, it is not one that we can hear. We get the benefits of equal temperament—a single well-defined frequency for every note—and the disadvantages are so small that we will never notice them.

Note	Cents difference
Fundamental	0
Major 2nd	3.91
Major 3rd	−13.69
Perfect 4th	− 1.96
Perfect 5th	1.96
Major 6th	−15.64
Major 7th	−11.73
Octave	0

Table 2.2: The number of musical cents by which the notes of the just intonation major scale differ from the corresponding notes of the equal temperament major scale. Positive numbers indicate that the just intonation note is higher, negative numbers that it is lower. The three highlighted entries have pitch differences large enough to be audible to the human ear.

The pitch difference is not the same for all notes of the scale, however. We can do the same calculation for the third note of the scale, for instance. In the just intonation scale this note has frequency $\frac{5}{4}f$ and in the equal temperament scale it is four half-steps up from the fundamental, giving it a frequency of $2^{4/12}f$. The equivalent calculation to Eq. (2.36) then gives

$$n = \frac{1200}{\log 2} \log \frac{\frac{5}{4}}{2^{4/12}} = -13.69 \text{ cents.} \tag{2.37}$$

This is now a larger pitch difference. 13.69 musical cents is enough that many listeners would be able to hear the difference between the notes. Many musicians state that they can hear that the equal temperament third is out of tune, although it is not something most of us register when we are listening to music under normal circumstances. Notice that the number of cents is negative in this case, meaning that the just intonation third is *lower* than the equal temperament third (by contrast with the second, which is higher).

We can continue these calculations for all notes of the major scale. Table 2.2 shows the results. As we can see, four notes—the second, fourth, and fifth of the scale, along with the octave—differ by amounts smaller than the ear can hear, but the other three—the third, sixth, and seventh (highlighted)—are off by larger intervals. For each of these three notes the just intonation scale has lower pitch than the equal temperament scale and the pitch differences are above the five cent threshold at which they would become apparent to the most sensitive ears. The pitch difference for the sixth in particular, which is over 15 musical cents, would be audible to most listeners. If you want to see for yourself, listen to the sound files `justscale.mp3` and `equalscale.mp3` to hear examples of the two major scales starting on middle C.

Files:
`justscale.mp3`
`equalscale.mp3`

When the equal temperament scale was first introduced in the 17th century, it provoked substantial opposition from musicians, many of whom felt that it was too far from the "ideal" pitch of just intonation and sounded out of tune. Equal temperament was opposed by some scientists and philosophers as well, who believed that the simple mathematical ratios of just intonation were a representation of a deeper, perhaps divine truth that should not be distorted. There was a long battle before superior convenience won the day and the equal temperament scale achieved the acceptance it has today.

For an entertaining account of the historical fight between supporters of just intonation and equal temperament tuning, take a look at the book *Temperament* by Stuart Isacoff.

Advanced material

2.4.3 Equal temperament and rational approximation

It is mathematically inevitable that just intonation and equal temperament do not agree perfectly, because the frequency multipliers for intervals in just intonation are all rational fractions, like $\frac{3}{2}$ and $\frac{4}{3}$ (see Fig. 2.8), whereas the multiplier $2^{1/12}$ used for the equal temperament scale is an irrational number. It is interesting to ask, however, how closely the rational fractions can approximate the equal temperament scale. If we are allowed to use any rational fraction, then we can approximate an irrational number as closely as we like, but the fractions used in just intonation are all "simple" fractions, meaning that the numbers in the numerator and denominator are quite small. We have fractions like $\frac{5}{4}$ or $\frac{9}{8}$, but not $\frac{1234}{5678}$. How good an approximation can we get with such simple fractions only?

One way to make this idea of simple fractions more formal is to consider the prime factors of the numerator and denominator. Every integer can be written as a unique product of prime factors and every fraction is a ratio of such products. For instance,

$$\frac{15}{8} = \frac{3 \times 5}{2 \times 2 \times 2}. \tag{2.38}$$

By definition there are no common prime factors on the top and the bottom of the fraction, since if there were they would cancel out.

One possible definition of a "simple" fraction is one that only contains a few prime factors—the fraction $\frac{15}{8}$ above, for instance, has five: one 3, one 5, and three 2s. Given that we must have different factors in the numerator and denominator, we have to use at least two different prime numbers, one on the top and one on the bottom. In the simplest case we could make fractions out of 2s and 3s only. Then every possible fraction would have the form $2^p 3^q$, where p and q are integers that can be positive, negative, or zero.

This approach is called *3-limit tuning*, because 3 is the largest prime number we are allowed to use. An example of a 3-limit scale is the so-called Pythagorean scale, a variant of just intonation that constructs a major scale using the following sequence of fractions:

$$\begin{aligned}
\text{Fundamental:}\quad & 1 = 2^0 3^0 \\
\text{2nd:}\quad & 9/8 = 2^{-2} 3^2 \\
\text{3rd:}\quad & 81/64 = 2^{-6} 3^4 \\
\text{4th:}\quad & 4/3 = 2^2 3^{-1} \\
\text{5th:}\quad & 3/2 = 2^{-1} 3^1 \\
\text{6th:}\quad & 27/16 = 2^{-4} 3^3 \\
\text{7th:}\quad & 243/128 = 2^{-7} 3^5 \\
\text{Octave:}\quad & 2 = 2^1 3^0
\end{aligned}$$

The more conventional just intonation scale uses *5-limit tuning*, meaning it makes use of the primes up to five—i.e., 2, 3, and 5—so that every fraction is of the form $2^p 3^q 5^r$, with the ratios being:

$$\begin{aligned}
\text{Fundamental:}\quad & 1 = 2^0 3^0 5^0 \\
\text{2nd:}\quad & 9/8 = 2^{-2} 3^2 5^0 \\
\text{3rd:}\quad & 5/4 = 2^{-2} 3^0 5^1 \\
\text{4th:}\quad & 4/3 = 2^2 3^{-1} 5^0 \\
\text{5th:}\quad & 3/2 = 2^{-1} 3^1 5^0 \\
\text{6th:}\quad & 5/3 = 2^0 3^{-1} 5^1 \\
\text{7th:}\quad & 15/8 = 2^{-3} 3^1 5^1 \\
\text{Octave:}\quad & 2 = 2^1 3^0 5^0
\end{aligned}$$

As long as the numbers p, q, and r are small integers then the fraction is "simple" in the sense proposed above.

To understand how much these scales differ from the equal temperament scale it helps to break them down into half-steps. In the equal temperament scale a half-step always corresponds to the irrational frequency multiplier $2^{1/12}$, exactly a 12th of an octave, but for scales that use rational numbers a half-step cannot be exactly a 12th of an octave, which necessarily means that some half-steps have to be larger than others if twelve of them are to add up to an octave.

Take the case of the Pythagorean scale. For this scale it turns out that half-steps come in two different sizes, known as the *minor* and *major Pythagorean semitones*. The minor semitone corresponds to the frequency multiplier

$$s_1 = \frac{2^8}{3^5} = \frac{256}{243} = 1.05350 = 90.2 \text{ cents}, \tag{2.39}$$

about 10 cents smaller than the full 100 cents of the equal temperament half-step. The major semitone is

$$s_2 = \frac{3^7}{2^{11}} = \frac{2187}{2048} = 1.6787 = 113.7 \text{ cents}, \tag{2.40}$$

about 14 cents larger than the equal temperament half-step. By combining these two semitones we can make all the notes of the Pythagorean scale. For instance,

$$s_1 s_2 = \frac{2^8}{3^5} \times \frac{3^7}{2^{11}} = \frac{3^2}{2^3} = \frac{9}{8}, \tag{2.41}$$

which gives us the second note of the scale, and

$$s_1^3 s_2^2 = \frac{2^{24}}{3^{15}} \times \frac{3^{14}}{2^{22}} = \frac{2^2}{3^1} = \frac{4}{3}. \tag{2.42}$$

And so forth. Every one of the fractions of the Pythagorean scale can be written in the form $s_1^{m_1} s_2^{m_2}$.

A convenient way to represent this is with the diagram shown in Fig. 2.13. The axes in this figure measure the numbers m_1 and m_2 of semitones and every possible combination of the two Pythagorean semitones is represented

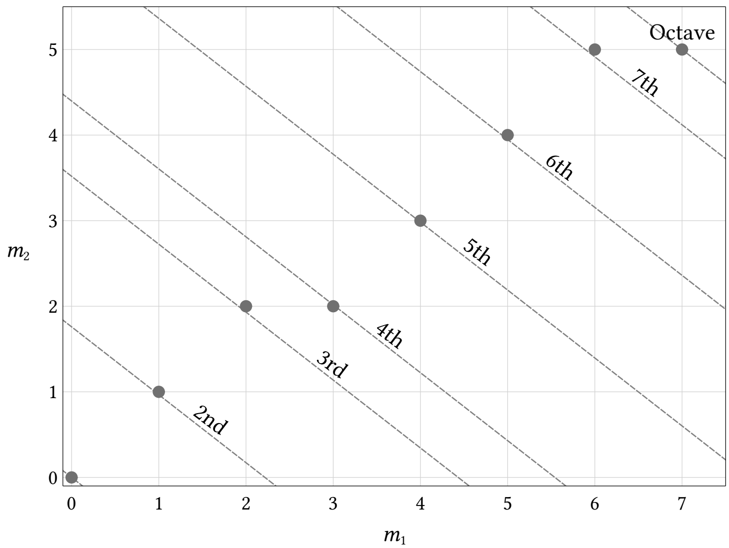

Figure 2.13: Visualization of the Pythagorean and equal temperament major scales. Each note of the Pythagorean scale can be reached by starting at the fundamental and raising the pitch by m_1 "minor semitones" and m_2 "major semitones," where m_1 and m_2 are integers. The dots represent the values of m_1 and m_2 for each of the eight notes of the Pythagorean major scale. The diagonal lines show where the dots would have to fall to agree with the equal temperament major scale. In all cases they are quite close to the lines, meaning the Pythagorean and equal temperament scales will sound closely similar, but small differences are visible.

by one of the grid-points in the figure. The eight highlighted points represent the eight notes of the Pythagorean major scale.

And what values of m_1 and m_2 correspond to the equal temperament scale? Integer values will not work since these give rational fractions, but if we allow non-integer values then, writing the frequency multiplier $2^{n/12}$ for the nth step of the equal temperament scale as $2^{n/12} = s_1^{m_1} s_2^{m_2}$ and taking logs of both sides, we get

$$m_1 \log s_1 + m_2 \log s_2 = \frac{\log 2}{12} n. \tag{2.43}$$

This is a linear equation in m_1 and m_2, meaning that there is a whole set of values m_1, m_2 lying on a straight line that all produce the interval of n equal temperament half-steps. The diagonal lines in Fig. 2.13 represent these sets, one for each of the eight notes of the major scale $n = 0$, 2, 4, 5, 7, 9, 11, and 12. For the Pythagorean scale to match the equal temperament scale exactly, the eight dots would have to lie on these eight lines. The extent to which they do not is the extent to which the Pythagorean scale differs from the equal temperament scale. As the figure shows, the agreement is actually very close. There is only the smallest of distances between the dots and lines, the largest gaps being for the 3rd and 7th notes of the scale.

Can we do the same kind of calculation for the just intonation scale? Indeed we can, but it is more complicated because just intonation is a 5-limit scale that uses the first three prime numbers and not just the first two like the Pythagorean scale. This means that there are three different semitones, given by

$$s_1 = \frac{5^2}{2^3 3^1} = \frac{25}{24} = 1.04167 = 70.6 \text{ cents}, \tag{2.44}$$

$$s_2 = \frac{3^3 5^1}{2^7} = \frac{135}{128} = 1.05469 = 92.9 \text{ cents}, \tag{2.45}$$

$$s_3 = \frac{2^4}{3^1 5^1} = \frac{16}{15} = 1.06667 = 111.7 \text{ cents}. \tag{2.46}$$

By appropriate integer combinations of these three semitones we can make all of the notes of the just intonation scale.

Unfortunately, this means that the diagram equivalent to Fig. 2.13 for the just intonation scale is three-dimensional, with axes representing the numbers m_1, m_2, and m_3 of each of the three semitones and two-dimensional planes (rather than lines) representing equal temperament. Such a three-dimensional diagram cannot be usefully reproduced on a two-dimensional page. Nonetheless, in principle at least, these methods can be used to represent how much each note of the just intonation scale differs from the corresponding note of the equal temperament scale (or *vice versa*).

Can we extend the same process to even higher dimensions by using more prime numbers? In principle the answer is yes, but in practice this is not done, at least in Western music. Three dimensions is as high as things go. When we multiply or divide the frequency of a note by factors of 2, 3, and 5 in 5-limit tuning we are shifting the pitch by intervals of an octave, an octave plus a fifth, and two octaves plus a third (see Fig. 2.7). The next largest prime factor that we might incorporate into this process would be 7 (which gives us *7-limit tuning*) but, as we will see in Section 4.3.4, multiplying the frequency by seven does not correspond to any standard musical interval in Western music. Thus factors of seven will not give us notes on the major scale, or any other conventional scale, and in practice sound rather out of tune. The same is also true of higher prime numbers—11, 13, 17, and so forth. The only ones that give us musically useful notes are 2, 3, and 5. Higher prime factors, including 7, are used in some other musical traditions and in microtonal and experimental electronic music styles. For instance, the Vietnamese string instrument called the dan bau uses frequency ratios based on factors up to 7, as do some bagpipes.

2.4.4 Meantone and other temperaments

We have seen two different versions of the major scale, based on just intonation and equal temperament. Just intonation gives us perfectly consonant musical intervals but requires us to retune our scale for each key we play in. Equal temperament allows us to use the same tuning for every key but gives imperfect intervals that are slightly out of tune.

These, however, are not the only possibilities. There are also a variety of other scales, or temperaments, that split the difference between just intonation and equal temperament in various ways, with intervals that are more consonant than equal temperament while still allowing us to play in more than one key. The best known is *meantone temperament*, which comes in several versions, varying in how much emphasis they place on the two competing priorities.

To construct a meantone scale, we start with a just intonation scale in a particular key, such as C major, then shift some of the notes part of the way toward their equal temperament counterparts to create a scale that is usable (if imperfect) in keys other than C major without retuning, but is simultaneously less out of tune than equal temperament. An instrument tuned to a meantone scale based on C major will still sound best for music in the key of C, but will typically also sound acceptable in keys that are harmonically close to C, meaning they share many of the same notes. As the key gets further away from C, however, the sound will get progressively worse.

For instance, the F major scale is close to C in the sense that it has all of the same notes as C major except for one (B gets replaced by B♭), so it will sound acceptable when played on an instrument with a C major meantone tuning. On the other hand, the F♯ major scale shares hardly any notes with C major (only B and F are common to both), so it will sound quite bad.

Related to meantone temperaments are the *circulating temperaments*, also called well temperaments, which take the idea a step further, shifting the notes further from their just intonation ideal to create a tuning where scales in every key are acceptably in tune. Circulating temperaments are in this respect similar to equal temperament. The crucial distinction is that different scales do still sound somewhat different in a circulating temperament, some being more in tune than others.

Meantone temperament, and particularly the version known as *quarter-comma meantone*, was widely used from the 16th century until the early 19th century (when it was superseded by equal temperament), and can still be heard on musical instruments such as pipe organs that date from earlier eras. An important feature of meantone (or circulating) tuning is that, while one can use it play in a range of different keys, the keys do not all sound the same. An instrument such as an organ would typically have been tuned in C major meantone and hence will sound most consonant for music in C, but more and more dissonant as one moves to progressively more distant keys. The varying character of the different keys was well understood by musicians of the time, and it was common for people to have favorite keys or to associate particular keys with specific moods or flavors. C major, for instance, was considered the purest and most perfect of keys, while D major was the key of celebration and triumph, as in Beethoven's *Ode to Joy* from the choral symphony. A♭ major on the other hand was allegedly the key of death and decay and B was the key of anger and despair.

Fanciful though these descriptions may seem, there is some truth behind them, given that the level of dissonance in the music depended on what key it was played in, and music of the 17th and 18th centuries deliberately exploited this dependence. A common compositional gambit was to start a piece of music in a consonant key close to the tuning key of the instrument (such as C major), then move through a series of key changes to more distant (and hence more dissonant) keys, thereby building tension, before moving back again to the starting key and resolving the tension at the end of the piece. This dramatic device is no longer available to us in the modern era of equal temperament tuning and music is arguably poorer for it, although equal temperament opens up many of its own musical possibilities that were not available in earlier times. At the very least, however, one should be aware that music written for meantone or circulating temperaments will not sound the same when played using modern tuning. It is not merely that the details of the tuning will be different: deliberate elements of tension and drama that derive from the consonance and dissonance of different keys will be absent. For this reason, historical temperaments still play a role in today's performance practice when playing music written with them in mind.

These considerations apply primarily to instruments that cannot adjust their tuning as they play, such as keyboard instruments, fretted strings such as lutes and guitars, and wind instruments, although skilled wind players can adjust tuning to some extent by controlling air pressure and lip tension. Other instruments, such as unfretted strings (violin, cello, and so forth) and especially the voice, can adapt their tuning as they play and hence need not stick to a single tuning for each note. The same note played or sung in different contexts can be tuned slightly differently. Even these instruments, however, are often played at the same time as less flexible instruments like organ or piano, and hence may be obliged to conform to a fixed temperament so that the ensemble is in tune with itself. An important exception is unaccompanied vocal music, including the substantial corpus of period liturgical music that survives from the era before equal temperament. In principle such music can be sung using any choice of tuning, including just intonation, meantone, or equal temperament, given sufficiently skilled performers.

2.4.5 Concert pitch and the frequency of A4

As described in Section 2.4, one can calculate the frequency of any note in the conventional equal-temperament 12-tone scale once one knows the frequency of any one starting note. From there, one just goes up or down the appropriate number of half-steps by multiplying or dividing by factors of $2^{1/12} = 1.05946$.

By convention, the starting note, the reference point from which the calculation begins, is the note A4, the A above middle C, which is defined to have a frequency of 440 Hz, known as *concert pitch.* Although this choice is the standard today, however, it has not always been so. The use of A4 as a reference point has been common for centuries, but the particular choice of 440 Hz is relatively recent, being formally adopted in the United States in 1936 and worldwide only in 1955. Before that there was no widely accepted standard and different choices were in use at different times and in different parts of the world.

In the Baroque era of classical music (approximately the years 1600 to 1750) pitches lower than 440 Hz were common. How do we know this? There are no recordings from that era—sound recording was not invented until the latter part of the 19th century—but we have physical evidence from historical musical instruments that have retained their pitch, such as pipe organs, and especially from tuning forks.

A tuning fork is a simple device used by musicians to produce a reliable note that you can tune an instrument to. A typical tuning fork is made of metal, about 15 cm

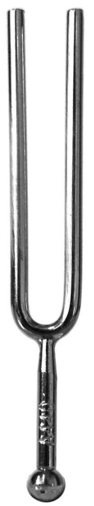

Figure 2.14: A tuning fork that, when struck, reliably produces a note with frequency 440 Hz.

long, with two prongs as shown in Fig. 2.14. You strike it on any convenient hard surface and the prongs vibrate and produce a note. Being made of metal, tuning forks are durable, and there exist a number of historical tuning forks that have come down to us over the centuries. We can use these to determine the pitch to which musicians tuned their instruments in years past.

For instance, a tuning fork that formerly belonged to the Baroque composer George Frederick Handel and which now resides in a museum in Ingolstadt, produces the note C5, but at a frequency of 512 Hz, somewhat lower than the 523 Hz we calculated in Eq. (2.26). If we go down three half-steps from C5 to A4, this means that Handel's A4 would have had a frequency of 422 Hz, about 38 musical cents lower than today's standard of 440 Hz, a difference large enough to be readily apparent to the ear. Because of the lower pitch standards in use in the Baroque era, music would have sounded somewhat different then and musicians playing early music today sometimes deliberately tune their instruments lower in an attempt to create an authentic sound.

Pitch standards rose following the Baroque era. Higher pitches can give music more energy and brilliance and there was a period of "pitch inflation" in the early 19th century when frequencies crept ever higher in pursuit of greater musical excitement. Some tuning forks from this period place A4 at 450 Hz or even higher. Musical excitement notwithstanding, this resulted in protests from instrumentalists and particularly singers, who strained to reach the highest notes, and in 1859 the government of France established a new and more moderate pitch standard of 435 Hz for A4, which was subsequently adopted by several other countries as well. Various attempts at further standardization took place over the next century, but it was not until 1955 that the International Standards Organization finally settled on the current value of 440 Hz for the frequency of A4. There are still some countries and organizations that use other pitch standards—442 and 443 Hz are both used by some orchestras and lower pitches are sometimes used when performing Baroque music as we have said—but 440 Hz is the only accepted international standard and it is the one we will use in this book.

The notion that higher pitches create more energy does still live on in some walks of musical life. In popular music, for instance, recording engineers discovered quite early on that slightly speeding up the tape of a recorded performance produces a brighter sound (as well as increasing the tempo, which can also raise the energy level). One cannot increase the speed very much without distorting the timbre of a performance—singing voices in particular can easily pick up a squeaky "chipmunk" sound if one is not careful. However, a small increase in speed, pushing the pitch up by no more than a half-step can, arguably, be beneficial. The practice was particularly common in the 1960s and 70s, and if you play along with recordings from that era, say on the piano, you may find that the recordings and the piano are not in tune with one another. An example is the song *Jumpin' Jack Flash* by the Rolling Stones, which is nominally in the key of B♭, but is about 40 musical cents above standard concert pitch because the recording has been sped up.

2.5 Musical notation

So far we have talked about specific musical notes and pitches using note names like "A4" and keys on the piano keyboard, as in Fig. 2.9. Musicians, however, have their own notation for notes, which are represented as ovals that sit on a set of horizontal lines, and this notation is a convenient way to succinctly describe musical notes, phrases, or complete compositions. It is not necessary to be fluent in musical notation in order to read and understand this book. Our primary focus is on the science of music, not on musical performance. But there will be times when we need to write down a snippet of music or a scale, and for those times it will be useful to have a basic grasp of musical notation and an ability to work out what it says. In this section we give a beginner's introduction to musical notation, focusing on the features that will be valuable when reading the rest of the book. If you can already read music, then you can safely skip this section.

2.5.1 The staff and notes

Music is written on a *staff*, a set of five horizontal lines that looks like this:

In British English the staff is called a "stave."

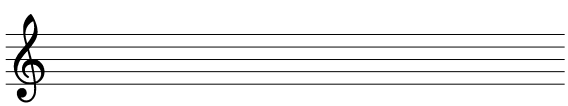

The curly symbol on the left is a *clef*. It is used to distinguish between several different conventions for what the lines mean. The particular clef shown here is the *treble clef* and it is the main one we will use in this book and the only one you need to be familiar with. Occasionally, to write very low notes, we will use the *bass clef*, which looks like this:

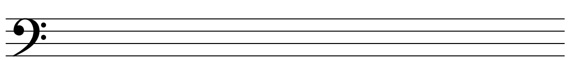

However, for the purposes of the book you will not actually need to be able to read notes written using the bass clef.

Notes are written as ovals, colloquially called "blobs," that fall either on the lines or between them, like this:

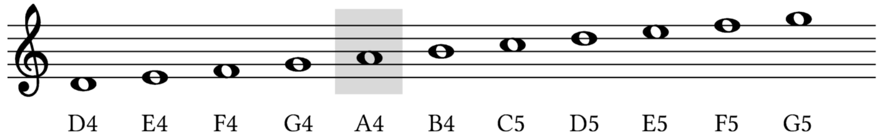

When performing music, notes are played in order from left to right, as one would read a book.

Lower blobs correspond to lower notes, higher blobs to higher ones, as indicated. Thus for instance the note highlighted in the gray box in the middle is A4, with frequency 440 Hz. The blobs to the right of that represent higher notes, going up the white keys of the piano—B4, C5, D5, and so forth—while those to left represent the notes going down—G4, F4, E4, and so forth. Thus the notes in this example range from D4 up to G5.

If we run off the top or bottom of the staff, we can add small extension lines, also called "ledger lines," to denote particularly high or low notes, like this:

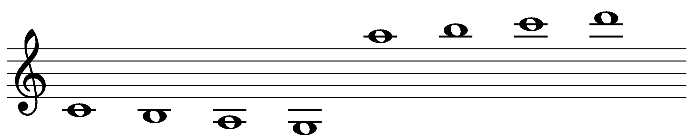

For instance, the first note in this example is the note C4, or middle C. A complete C major scale would be written like this:

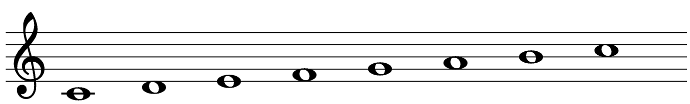

The "black keys" on the piano, the sharp and flat notes, are represented by placing a sharp or flat sign *before* the note in question (not after it), like this:

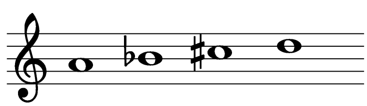

This represents the notes A4, B♭4, C♯5, and D5. In cases where sharps or flats are used many times in the same piece of music an alternative and convenient convention is to collect the regularly used ones into a *key signature* at the start of the line like this:

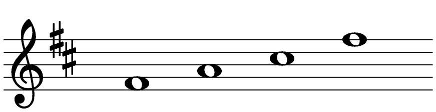

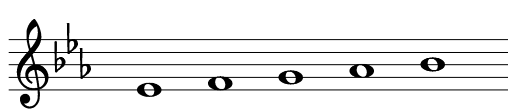

In the example on the left there are two sharp signs in the key signature, one on the note F and one on C. This means that *all* Fs and Cs are actually to be played as F♯ and C♯, wherever they occur in the music. That includes Fs and Cs in any octave, not just the particular notes F5 and C5 that the sharp symbols are drawn on. So, for instance, the F4 in the example on the left—the first note of the four shown—would be played as F♯4. In the example on the right the key signature contains three flats, on B, E, and A, telling us that every occurrence of any of these notes, in any octave, is to be played as B♭, E♭, or A♭.

2.5.2 Scales

We have looked at the major scale in some detail in this chapter. Written in musical notation, starting on C4, it looks like this:

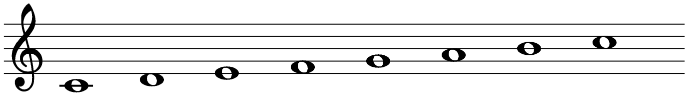

Although it is common to write scales in ascending order of their notes like this, they can be played in any order—ascending, descending, or some random order. The point of a scale is that it gives us a set of notes that sound pleasing together. It is the collection of notes that matters, not the order in which they are played.

The major scale is not the only scale in common use. There are many others, each consisting of some selection of notes from the 12 tones of the equal temperament system. As an example of the use of musical notation, let us take a look at some of the other scales that are common in Western music. The next most common scale after the major scale is the *minor scale*:

Technically this is the *melodic minor* or aeolian scale. There is also the *harmonic minor* scale, which is slightly different.

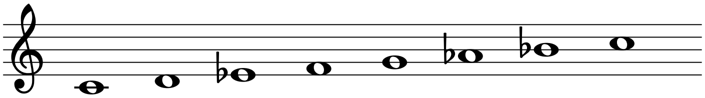

Music using a minor scale is often considered somber or melancholy in sound. Examples include the funeral march from Chopin's 2nd Piano Sonata and the traditional song *House of the Rising Sun.*

The *pentatonic scale* is a scale with only five notes, or six if you count the note repeated at the top:

This is the *major pentatonic* scale. You might occasionally also encounter the less widely used *minor pentatonic* scale.

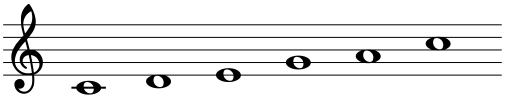

Pentatonic scales are common in folk music of various traditions and also occur widely in popular music styles. And example of the use of the pentatonic scale is the hymn *Amazing Grace*. The pentatonic scale is also the primary scale in many types of traditional Chinese music, and some instruments, such as the guzheng, are tuned so as to play only notes of the pentatonic scale.

The *blues scale* is a scale used in blues, jazz, and related musical styles:

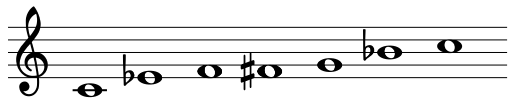

Using the blues scale can impart an instant "jazzy" sound to a piece of music. An example of a song written entirely using the blues scale is *Moanin'* by Bobby Timmons, famously recorded by Art Blakey's Jazz Messengers.

The *chromatic scale* is a scale composed of all 12 notes of the 12-tone scale:

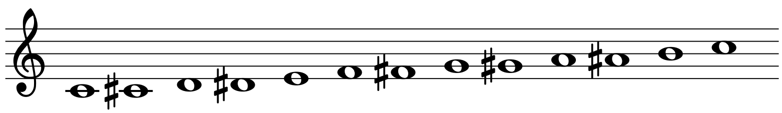

The chromatic scale is not widely used in melodic writing, but formed the basis for the 12-tone serialist style of early 20th century composers like Arnold Schoenberg and Igor Stravinsky, and is sometimes used in jazz or classical music for effect. An example is *The Flight of the Bumblebee* by Nikolai Rimsky-Korsakov.

These are only a small fraction of the many possible scales that can be created by putting together combinations of notes. There are probably hundreds of scales that have been created at one time or another, each with its own particular musical aura.

		Name	
Note	Length	US	European
𝅝	1	Whole note	Semibreve
𝅗𝅥	$\frac{1}{2}$	Half note	Minim
♩	$\frac{1}{4}$	Quarter note	Crotchet
♪	$\frac{1}{8}$	Eighth note	Quaver
𝅘𝅥𝅯	$\frac{1}{16}$	Sixteenth note	Semiquaver

Table 2.3: Notation and names for notes of various lengths.

2.5.3 Note duration

Musical notes can have different durations. Some are long and drawn out, others are short and rapid. Durations of notes are indicated in musical notation by different styles of blobs. The most commonly used durations are based around the mathematics of powers of two. The blobs in all the examples above are *whole notes*, which are usually the longest notes in typical music. A whole note is also sometimes called a *semibreve*, particularly in classical and European musical traditions. Exactly how long a whole note will be depends on how fast the music is played, but a typical duration would be around two seconds.

Other common lengths of notes are half notes, quarter notes, eighth notes, and sixteenth notes, which, as their names suggest, have lengths that are a half, a quarter, an eighth, and a sixteenth of the length of a whole note. In classical and European traditions these are also called *minims*, *crotchets*, *quavers*, and *semiquavers*, respectively. The European note names, however, have no particular logical structure and for those not used them are hard to remember, so we will not use them in this book.

Table 2.3 shows the styles of blobs used to represent each of these durations. For instance, a half note is written as an open blob with a vertical stem attached. A quarter note is a solid blob with a stem attached. An eighth note is a solid blob and a stem with a "flag" or "tail" on the end. A sixteenth note is the same but with two tails. When several eighth or sixteenth notes are written in a row, their tails may be joined together into a continuous bar or "beam." Here is an accelerating passage of music that starts with long slow notes and progresses to shorter faster ones:

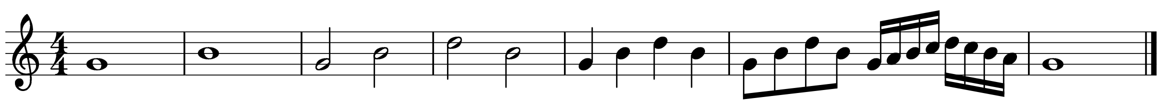

This example shows a couple of other features as well. First, notice the vertical lines drawn across the staff at various points. These are *bar lines* and they divide the music in *bars* or *measures*, which are successive portions of music that each last for the same amount of time. In this example, for instance, the first two measures contain one whole note each and the next two contain two half notes each—and hence last the same length of time, since two halves make a whole. Similarly the next measure contains four quarter notes, again adding up to the same total amount of time, and so forth.

Bar lines are not strictly necessary for writing music. One could write music without them and it would sound the same when played. But measures are useful because they help the musician keep track of the passage of time by providing a visual timescale for the musical staff. They also have musical significance in that music normally has a natural rhythm to it, a steady pulse or beat that contributes to the feel of the music, and beats are usually grouped together into sets, often (again) according to powers of two. The most common example is music "in four," meaning it naturally divides into groups of four beats. A large majority of popular music and Western classical music is in four. The musical measures denoted by the bar lines in written music are normally chosen to coincide with the natural groups of beats. Again this makes the music easier for the musician to follow.

To alert the musician to the particular grouping of beats in use, written music normally starts with a *time signature*—the numbers at the start of the line in the example above. In this case the numbers are $\frac{4}{4}$, which means each bar consists of four beats, each of which is a quarter note long. In general, the upper number indicates the number of beats and the lower one indicates the length of each beat, with a 4 indicating a quarter note length, an 8 indicating an eighth note, and so forth. If you want to hear an example of music in $\frac{4}{4}$ time, just turn on the radio. Almost anything you hear will be in four.

The common use of the quarter note as the beat does have the result that a "whole note" and "one beat" are not the same thing: there are four quarter-note beats in a whole note. This can be confusing if you are encountering it for the first time, in which case it may be worth taking a moment to review and understand the distinction. Whole notes, half notes, quarter notes, and so forth are merely conventional names for different lengths of notes, but any of them can play the role of the beat in a piece of music. The lower number of the time signature tells you what the length of a beat is in any specific case.

After $\frac{4}{4}$, the next most common time signature is $\frac{3}{4}$, which means three quarter-note beats to a bar, also called "waltz time" because of its use in the ballroom dance. Think of the *Blue Danube Waltz* by Johann Strauss, *Edelweiss* from *The Sound of Music*, or Kermit's song *The Rainbow Connection* from *The Muppet Movie*. The $\frac{3}{4}$ and $\frac{4}{4}$ time signatures account for most of the music one hears, but others, including such oddities as five-, seven-, and thirteen-beat time signatures, are found occasionally. The theme from *Mission Impossible* is in $\frac{5}{4}$ time, as is *Mars* from *The Planets* by Gustav Holst. Pink Floyd's song *Money* is in $\frac{7}{4}$ and so is the verse (but not the chorus) of *All You Need Is Love* by the Beatles, and $\frac{13}{8}$ is used for parts of the song *Skimbleshanks the Railway Cat* from the musical *Cats*. One may also encounter *compound* time signatures, which consist of "major" beats subdivided into less prominent "minor" beats. The most common example is $\frac{6}{8}$ time, in which a measure has two major beats, each divided into three minor ones, for a total of six minor beats in the whole measure. Familiar examples include *What a Wonderful World* by Louis Armstrong and *America* from the musical *West Side Story* (which alternates between $\frac{6}{8}$ and $\frac{3}{4}$ time). Another common compound time signature is $\frac{12}{8}$, in which a measure has four beats with three subbeats each. $\frac{12}{8}$ found use particularly in the 1950s and 60s in popular ballads such as *Natural Woman* by Aretha Franklin and *Can't Help Falling in Love* by Elvis Presley.

Chapter summary:

- Musical notes correspond to sounds with **periodic waveforms** and the **frequency** of the waveform—the number of cycles per second—determines the pitch of the note.
- The human ear can hear frequencies from about 20 to 20 000 cycles per second—also called hertz and denoted "Hz"—but musical notes only occupy the lower end of this range from about 20 to 4000 Hz.

- A musical interval—the change in pitch between one note and another—is created by **multiplying** the frequency by some number. For instance, to go up an octave you multiply the frequency by two.
- When two notes with different frequencies are played at the same time they sound pleasant together, or "consonant," if one frequency is a **simple multiple** of the other, such as 2 times, or 3 times, or $\frac{3}{2}$ times. We can use this principle to assemble a **scale**, which is a collection of notes that sound pleasing together. The most common scales in Western music are the major and minor scales, although there are many others as well.
- In practice, scales defined in this way using simple multipliers are inconvenient and so modern music is based instead on dividing the octave into twelve equal parts called **half-steps**. To go up one half-step, you multiply the frequency by $2^{1/12} = 1.05946$. To go down a half-step you divide by the same amount. This gives you a good but not perfect approximation to the notes of the major scale, along with five other notes, the "sharps and flats," represented by the black keys on the piano keyboard.
- Scales using frequencies defined in this way are called **equal temperament** scales, in contrast with those using the simple frequency ratios, which are called **just intonation** scales.
- Notes are identified by letters A to G, plus a number to indicate which octave they fall in. Sharps and flats are indicated by a sharp symbol ♯ or flat symbol ♭ after the letter name. The note "middle C," for instance, is C4 and the note a half-step above it is C♯4. The notes of the piano span just over seven octaves from A0 to C8.
- By international standard, the frequency of the note A4 is fixed at **440 Hz**. From this, we can calculate the frequency of any higher note using the formula

  $$f = 440 \times 2^{n/12},$$

  where n is the number of half-steps from A4 to the note in question. Similarly for notes lower than A4 we have

  $$f = \frac{440}{2^{n/12}}.$$

- Pitch intervals smaller than a half-step are measured in **musical cents**. One musical cent is a hundredth of a half-step and corresponds to multiplying or dividing the frequency by $2^{1/1200} = 1.000578$. The number n of cents between two frequencies f_1 and f_2 is given by

  $$n = \frac{1200}{\log 2} \log \frac{f_1}{f_2}.$$

- A basic knowledge of **musical notation** will be useful for reading the remainder of this book. A brief introduction can be found in Section 2.5.

Exercises

2.1 The waveform of a certain sound looks like this:

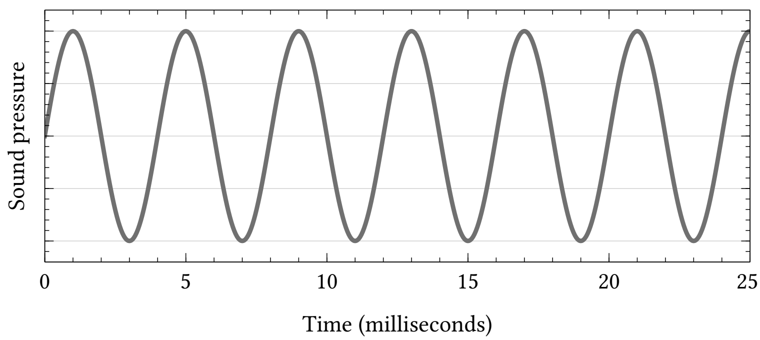

a) Give an estimate of the frequency of the sound.
b) What is the period in milliseconds?
c) What is the wavelength in meters?

2.2 The lowest frequency the human ear can hear is about 20 Hz and the highest is about 20 000 Hz. What are the corresponding wavelengths?

2.3 A certain note has a frequency of 440 Hz. Separately for just intonation and for equal temperament, calculate the frequencies of the notes that are

a) Up an octave from this
b) Down an octave
c) Up a major third
d) Up two fifths

2.4 What is the musical interval between the following pairs of frequencies?

a) 200 Hz and 800 Hz
b) 400 Hz and 600 Hz
c) 600 Hz and 1500 Hz

2.5 By what factor would you multiply a frequency to raise the pitch by these intervals in just intonation?

a) A major sixth
b) One octave plus a major second

c) Two octaves plus a major third

2.6 A certain note has a frequency of 220 Hz. What is the frequency of the note five half-steps higher, in equal temperament tuning?

2.7 How many half-steps are there between two notes with frequencies 392 Hz and 587.3 Hz?

2.8 The lowest note on a standard 88-key piano is A0, but some pianos go as low as F0. What is the frequency of F0?

2.9 Suppose we number the keys of the piano from 1 to 88, going from left to right. Derive a general formula for the frequency f_n of key number n.

2.10 Consider a major scale starting on the note A4 at 440 Hz and going upwards.

a) Calculate the frequency of the major third (the third note of the scale, which is C♯5 in this case) if the scale is played using just intonation and if it is played using equal temperament.
b) Is the difference between these two notes big enough for a human listener to hear?

2.11 Consider this note:

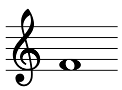

a) What is the name of this note? Give the full name, including the letter name and the octave number.
b) What is the frequency of the note in hertz to two decimal places (in equal temperament and at standard pitch)?
c) What is the wavelength in meters?
d) Another note has $\frac{1}{3}$ the wavelength. What note is it?

2.12 Most orchestras tune to standard concert pitch in which A4 has frequency 440 Hz, but a few, such as the New York Philharmonic Orchestra, use a higher frequency of 442 Hz.

a) How many musical cents are there between these two frequencies?
b) Is this difference big enough for a human listener to hear?

2.13 Consider this musical passage:

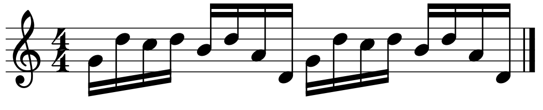

a) What does it mean that the music has a "$\frac{4}{4}$" time signature?
b) If this passage is played at 120 beats per minute, how long, in seconds, is each note?