Hearing math and seeing music: a workshop on pitch perception and temperament

Evelyn Lamb
Dept. of Mathematics, University of Utah
155 S 1400 E, Salt Lake City, UT 84112
lamb@math.utah.edu

This worksheet can be found online at http://www.evelynjlamb.com/bridges2015/.

Audacity Crash Course

1. If you haven't yet, download Audacity: http://audacityteam.org/

2. Open it.

3. On the Tracks menu, select Add New and then Audio Track.

4. On the Generate menu, select Tone. Choose your waveform, frequency, amplitude, and duration and click OK. You should now have a track with some kind of sound wave in it. Go you!

5. You'll be creating a lot of tracks today, so it's a good idea to name them as you go along. You'll probably want to name them by frequency and waveform.

6. As you add more tracks, note that there is a "mute" button on each track so you can listen to only a few at a time.

7. You can zoom in on a track to see the waveform by going to View and then Zoom In.

8. If you have several tracks going, zooming in won't help you visualize the combination of the wave-forms. If you would like to visualize that, go to the File Menu and click Export Audio or Export Selected Audio. Name the file whatever you want. Then reopen the file. It will be just one track, and when you zoom in, you can see the waveform. Another way to visualize these waves is to graph them on https://www.desmos.com/ or your online graphing calculator of choice. For example, the function $y = \sin(440x)$ would represent an A 440. If you go this route, it will probably be more convenient for you to reduce the coefficients in the sine but keep whatever ratios you have the same. As you go along, you might want to stop every once in a while and look at the music on Audacity, Desmos, or other graphic calculator.

Figure 1 : *A piano keyboard for reference. Perfect fifths are 7 half steps apart. For example, C and G. Octaves are 12 half steps apart.*

Some handy music terms

- Sharp: $N\sharp$ is the note one half step higher than N (to the right on the keyboard).

- Flat: $N\flat$ is the note one half step lower than N (to the left on the keyboard).

- Enharmonic equivalent: a note that is the same on a piano keyboard but named differently. For example, $E\flat$ and $D\sharp$. Today, we usually treat those notes as the same because they are the same in equal temperament, but some tuning systems will have differences between enharmonic equivalents.

- Intervals, based on number of half steps or piano keys between two notes:

  - 0: unison
  - 1: half step
  - 2: whole step
  - 3: minor third
  - 4: major third
  - 5: perfect fourth
  - 6: augmented fourth, diminished fifth, or tritone
  - 7: perfect fifth
  - 8: minor sixth
  - 9: major sixth
  - 10: minor seventh
  - 11: major seventh
  - 12: octave

- Chord: two or more pitches sounded together.

- Triad: Three notes played together. A major triad (e.g. C-E-G) has a major third and perfect fifth from the lowest note. A minor triad (e.g. D-F-A) has a minor third and perfect fifth from the lowest note.

- Timbre: sound quality. A flute sounds different from a violin playing the same pitch or a voice singing it. Timbre is the word we use to talk about those differences.

Intervals, well- and ill-tempered

Pythagorean tuning

1. Use Audacity to generate some square or sawtooth waves with frequency ratios of $2:1$ and $3:2$. We interpret those frequency ratios as octaves and perfect fifths, respectively. (Square and sawtooth waves are not beautiful, but sine waves cause some perceptual problems–see "how to trick yourself with sounds" below. I find the "square, no alias" tone to be the least annoying of the non-sine options on Audacity.)

We are used to an octave having 12 chromatic pitches in it, and if you took some music classes at some point, you might be familiar with the "circle of fifths." If you keep skipping 7 half-steps on a piano, you'll play a circle of fifths. The circle of fifths starting on C is C-G-D-A-E-B-F#-C#-G#-D#-A#-F-C.

The 3 : 2 and 2 : 1 ratios give us a way to find the frequencies of the 12 notes of the scale starting with one frequency. We'll call tuning using just these intervals Pythagorean tuning.

2. Starting with any frequency you want (A=440 or 220 is canonical, but you can use anything you'd like), use Audacity to generate square or sawtooth waves for some pitches in the chromatic scale with Pythagorean tuning. Because no power of 3 is equal to a power of 2, there will be some problems. Your usable intervals will be most convenient if you build "out" from your starting note rather than just "up". That is, if you start on A, find E by multiplying by 3/2 and D by multiplying by 2/3, then keep working from there. You can multiply or divide by 2 as necessary to keep pitches in the right octave.

I've written a Sage function that will spit out the frequencies of the notes in a chromatic Pythagorean scale. You can find it (read-only, so you'll have to copy it to your computer or the Sage math cloud) at `http://bit.ly/1erLkEL`. Pythag(n) spits out the frequencies for a scale centered on the frequency n. The function Pythag1(n) spits out the frequencies building up from the frequency n. The Sage worksheet also has an equal temperament function ET(n,m).

3. Pythagorean tuning makes beautiful fifths, but it has some problems. Play a major chord that has been generated with Pythagorean tuning. Does it sound good? Play the note that you started with and the note that you got by coming "back" to it using the circle of fifths. Do they sound like octaves? The difference between this note and the true perfect octave is called the Pythagorean comma. When you round that down to a true octave, you can hear the wolf fifth. In the Pythagorean scale centered on A, that will be the fifth between E♭ and G♯. In the Pythagorean scale building up from A, it will be the fifth between A and D. (Tip: sometimes intervals are easier to hear in one octave than another. That's why I've included the scales for both A=220 and A=440.)

Just intonation, meantone, etc.

4. We hear the frequency ratio of 5:4 as a "perfect" major third. What is the frequency ratio we are actually getting when we generate a C and E using Pythagorean tuning? Use Audacity to generate a major chord that has a perfect major third and a perfect fifth in it. (We call this just tuning.) Compare the sound of this chord to the Pythagorean major triad you got earlier. The frequency ratio for a just minor third is 6:5.

5. Just intonation is an umbrella term for tuning systems based perfect fifths and thirds. It's a game of compromises. Experiment with some just intonation scales. If you create a scale that plays some chords beautifully, it will probably sound very bad playing other chords. Play around with this.

6. In meantone tunings, perfect fifths are all "shaved" just a bit to give better thirds. The most common meantone tuning was quarter-comma meantone. In that tuning, a stack of fourth fifths was tempered to be exactly two octaves and a perfect major third. (That is, the frequency ratio of fourth fifths to the starting note was 5:1.) Experiment with quarter-comma meantone tuning. There is a wolf fifth in this tuning as well, this time one that is wider than a perfect fifth. There are an awful lot of other tuning systems based on similar ideas. You can probably come up with your own to play with if you want, or you can check out the technical library at this harpsichord page: `http://www.hpschd.nu/index.html?nav/nav-4.html&t/welcome.html&http://www.hpschd.nu/tech/tmp/mean.html`. (This link goes to the quarter-comma meantone tuning page. You can find other temperaments in a drop-down menu at the bottom of that page.)

Equal temperament

7. Tuning systems, or temperament systems, have always required compromise. The compromise we currently use most of the time is called equal temperament. Fifths are shaved uniformly so they all sound the

same and 12 of them close up perfectly to an octave. Given the fact that we hear frequency ratios of equal size as equal steps in music, this means the frequency ratio of a half step on the keyboard is $2^{1/12} : 1$.

8. Use that frequency ratio to figure out the size of a "perfect" fifth on a piano and compare it to a Pythagorean perfect fifth. Generate equal tempered major and minor triads. Compare these to other major and minor triads you've made.

9. What do you hear when you play the perfect and equal tempered fifths together? You can see the reasons for the "beats" you hear by looking at a graph of the waves. (Play around with sine waves on this one as well.)

10. Which intervals are well-approximated in 12-tone equal temperament? Which intervals are considerably worse than their "pure" versions?

Breaking the 12 barrier

All of our explorations until now have relied on the fact that we are trying to fit 12 pitch classes into an octave. What happens if we try to fit a different number?

11. An obvious way to increase the number of pitches available is to make a 24-tone equal tempered octave, adding quarter tones to our standard equal temperament. Play with a 24-tone scale. Do we get better approximations of any intervals?

12. A 19-tone equal tempered scale has been used for experimental tunings. Compute the frequencies of notes in a 19-tone scale. What common intervals are well-approximated in the scale? Which ones are not as well approximated as in 12-tone equal temperament? The Sage worksheet has functions to find frequencies and ratios for equal temperament scales with any number of pitches in the octave: `http://bit.ly/1erLkEL`.

13. One reason 12 works well for the number of chromatic pitches in a scale is that $7/12$ is very close to $\log_2(3/2)$. In fact, it is one of the convergents in the continued fraction for $\log_2(3/2)$. Two other convergents are $24/41$ and $31/53$. Play around with those scales and find out which intervals are well approximated there.

14. If you care about intervals other than perfect fifths, you can use the continued fraction representation of the just interval you're aiming for to find equal temperament scales that will approximate your interval well. The function CF(a,n) will give you the first n convergents in the continued fraction for $\log_2(a)$.

Seeing the music

Everything that we did can be visualized with sine waves on a graphic calculator, and I hope you've played with it as you've gone along. (Note that we made square or sawtooth waves instead of sine waves.) Sound waves are additive, so we can visualize a chord as the addition of sine waves with different frequencies. I find it especially interesting to visualize the "beat" effect we heard in number 9.

15. Sawtooth and square waves are infinite sums of sine waves with frequencies that have integer ratios. Play around with adding sine waves with different frequencies and amplitudes together and see if you can start to get things that look like sawtooth and square waves.

How to trick yourself with sounds

Sounds that come from voices and real instruments are not perfect sine (or sawtooth or square) waves. They are made up of a "fundamental" sine wave and waves that have integer frequency ratios to that fundamental. (Of course, that is an oversimplification as well, but we do what we can.) We call these harmonics or overtones.

16. In a new file in Audacity, create sine waves that represent a fundamental pitch and some of its overtones. You can change the timbre, or sound quality, by changing which overtones are included and their amplitudes, which corresponds to how loud each overtone is.

17. When you have made a fundamental and some of its overtones, click "mute" on some of the tracks so that only the even harmonics are playing. What do you hear? Add any odd harmonic. Now what do you hear?

18. Graph the effect you just heard, which is called the "missing fundamental." Fun fact: your telephone uses the missing fundamental effect because it can't pick up frequencies that are as low as the fundamental frequencies of most adults' voices.

19. Explore the limits of the missing fundamental. How few different frequencies can you include and still hear it? What happens if your frequencies are close but not exactly in a missing fundamental relationship?

Some researchers have been experimenting with changing the frequencies and amplitudes in tones and seeing whether people hear the missing fundamental or not. You can test yourself here: `http://www.khoomei.com/forum/uploads/1180135123.mp3` or by following the mp3 link here: `http://khoomei.com/forum/viewtopic.php?f=4&t=757`. You will hear pairs of notes and have to say whether it sounds like they go up or down. The second link has a "key" partway down the page.