

# The Magical Mathematics of Music

by

Jeffrey S. Rosenthal

*(Dr. Rosenthal is a professor in the Department of Statistics at the University of Toronto, and is an amateur musical performer who plays several instruments. His book "Struck by Lightning: The Curious World of Probabilities" is being published in September, 2005 by HarperCollins Canada.)*


The astronomer Galileo Galilei observed in 1623 that the entire universe "is written in the language of mathematics", and indeed it is remarkable the extent to which science and society are governed by mathematical ideas. It is perhaps even more surprising that music, with all its passion and emotion, is also based upon mathematical relationships. Such musical notions as octaves, chords, scales, and keys can all be demystified and understood logically using simple mathematics.

## **PITCH: WAVE FREQUENCIES**

Music appears to be transmitted by magic, escaping from your expensive stereo – or a loudly passing car radio, or a guitar-strumming maestro – and accosting your eardrums in one fell swoop. In fact, sound progresses as a wave through the *air*, and sound cannot be produced without an atmosphere. (Or, as the horror movies would say: in space no one can hear you scream.)


A sound wave creates minute pockets of higher and lower air pressure, and all the sounds we hear are caused by these pressure changes. With music, the frequency at which these pockets strike your ear controls the pitch that you hear.

For example, consider the note called "Middle C" (usually the first note learned in piano lessons). This note has a frequency of about 262 Hertz. That means that when Middle C is played, 262 pockets of higher air pressure pound against your ear each second. Equivalently, the pockets of air arrive so quickly that one pocket strikes your ear every 0.00382 seconds. We can draw a graph by putting an X at every time when a pocket of air arrives:



This graph provides a sort of “picture” of Middle C. By itself, it does not tell us much. However, such graphs provide a new perspective on the *relationships* between different musical notes.

A basic rule is that higher-pitched notes have a higher frequency, corresponding to *more frequent* air pocket arrivals. For example, the note Middle G (seven semi-tones *higher* than Middle C) has a frequency of about 392 Hertz, corresponding to 392 air pockets per second, or a time period of 0.00255 of a second between arrivals:



With the *higher* note (Middle G), the air pockets arrive *more frequently* – corresponding to a higher frequency, and thus to more X’s in the graph.

If you listen carefully to an ambulance siren or a train whistle, you will notice that the noise sounds *higher* while the vehicle is approaching, and *lower* after the vehicle has passed by. This is because the approaching movement compresses the X’s together, making them arrive more frequently and produce a higher pitch, while the departing movement stretches out the X’s and produces a lower pitch. This is musical frequency in action.


So how does this help us? Well, by using knowledge of sound frequencies carefully, such musical mysteries as octaves and chords can be unraveled.

## DOUBLE THE FUN: OCTAVES

No musical notes fit together better than those which are exactly one octave apart. Pairs of notes like Middle C and High C. Or Middle G and High G. Such pairs arise constantly in popular music, as in the first two notes of the classic *Somewhere Over the Rainbow*, or the initial “I’m sing. . .” of *Singing in the Rain* – or the first two notes of the third line of *Happy Birthday*.

Often in a symphony, one instrument (like a violin) plays a note one octave above that of another instrument (like a double bass). Piano arrangements inevitably include the same note at different octaves. When improvising in blues or jazz music, it hardly matters which *octave* you choose to play, just which notes *within* the octave you choose to play. Notes which are one octave apart sound almost like the same note, or at least a new variation on the same note.

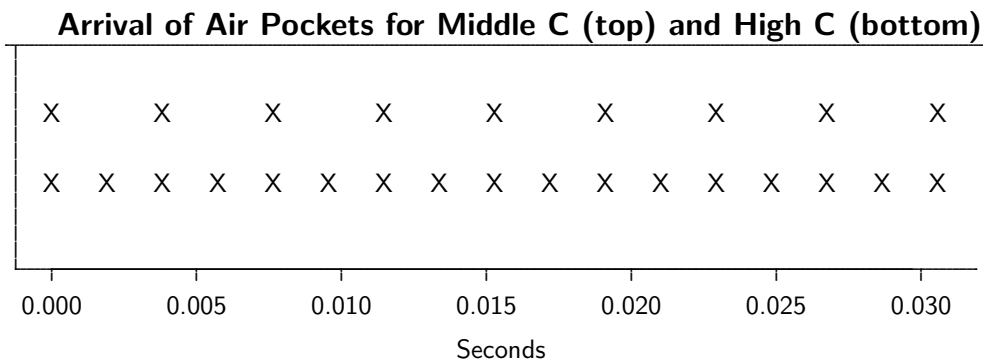
What is it that makes notes one octave apart fit together so well? Consider High C, the note which is exactly one octave above Middle C.



For High C, the air pockets arrive once every 0.00191 seconds. This is precisely *half* the time period (or, twice the frequency) of Middle C. The arrivals of High C are not random, or arbitrary, or picked out of a hat. Rather, they are precisely *twice as frequent* as the arrivals for Middle C.

This is a universal truth of octaves. Going up one octave is always the same as *doubling* the frequency. Just as the frequency of High C is twice that of Middle C, so the frequency of High D is twice that of Middle D, and so on. Notes one octave apart are always bound together by this basic relationship.

If we play both Middle C and High C at the same time, then their air pockets fit together in a certain pattern:



Since the frequency of High C is exactly twice that of Middle C, *the two notes line up perfectly*. Every *two* air pocket arrivals for High C correspond perfectly to *one* arrival for Middle C. The X's for the two notes match up exactly. That is why the two notes fit together so well, and indeed sound almost the same. Their high pressure air pockets arrive in perfect synchronisation.

Similarly, that is why High D fits in so well with Middle D, and indeed why any note fits in so well with the note one octave above or below. In every case, the frequency of one note is exactly twice the frequency of the other, leading to a perfect fit.

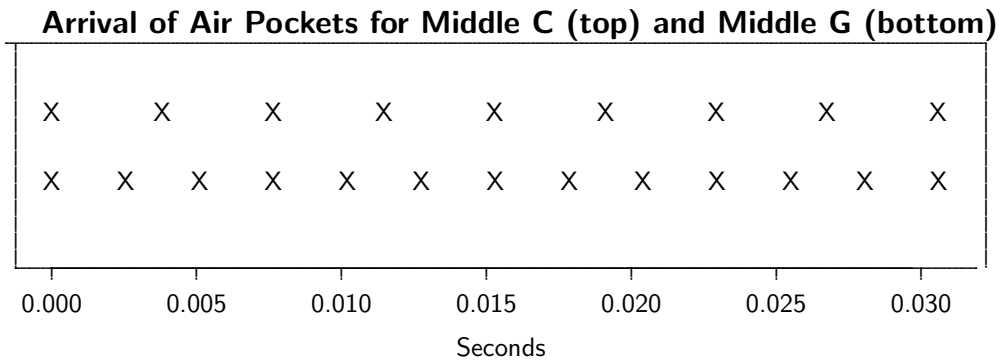
One easy way to see this is on a guitar. There, pressing a guitar string at the twelfth fret cuts the string precisely in *half*. This makes it vibrate precisely *twice* as quickly. This leads to a frequency which is twice as high, and thus corresponds to a note one octave higher. For example, the first string of a guitar is normally tuned to Middle E. Playing that same string with the twelfth fret pressed produces a High E – a note one octave higher, with frequency twice as high.

The same pattern continues as we increase the octaves. The note two octaves above Middle C (sometimes called High High C) has a frequency *four times* that of Middle C. And the note an octave above *that* one, has a frequency *eight times* that of Middle C. In principle, we could keep increasing the octaves, and doubling the frequencies, forever – but after a certain point, the notes would be so high that only dogs could hear them.

## IN PERFECT HARMONY

Notes an octave apart do indeed fit together well. In a sense, they fit together *too* well. If you play Middle C and High C together, then there is no discordance at all. However, the result sounds sort of *hollow*, or even *boring*. More interesting combinations like *chords* result if we bring in other notes, too. But which notes sound good together, and which ones don't?

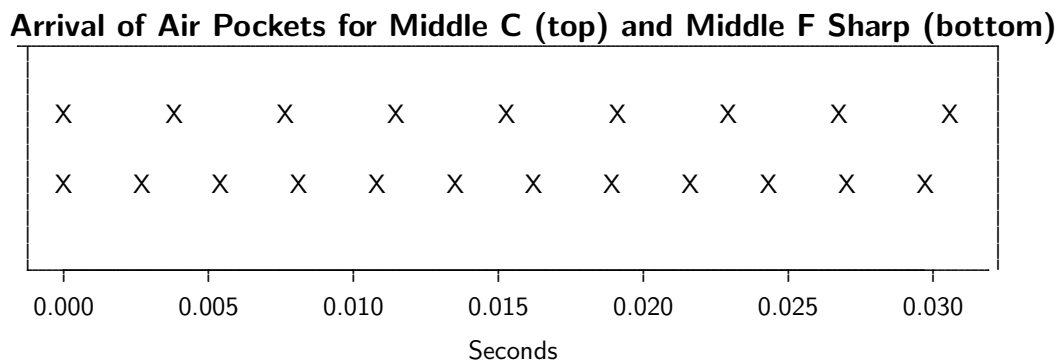
Consider the notes Middle C and Middle G. In fact, these notes form a “fifth” interval and fit together well. But why is this? Well, let’s examine their wave graphs together:



We see that every *second* arrival for Middle C, lines up almost perfectly with every *third* arrival for Middle G. Once again, *the two wave patterns fit well together*. The two notes complement each other, rather than clashing. That is why lots of music involves the notes C and G – from the first two notes of the chorus of the traditional song *Land of the Silver Birch*, to the first two notes of the *Star Wars* musical theme.

Since 5 time periods for Middle G correspond to 4 time periods for Middle C, we can also say that the frequency of Middle G is about  $5/4$  times as large as that of Middle C. Since  $5/4$  is a nice, simple fraction, that is why C and G sound good together.

By contrast, the notes C and F-Sharp do not fit together well. If those two notes are combined, they don’t sound quite right. They rarely appear as a pair in musical compositions or songs of any kind, at least not in a sustained way. A look at their graphs explains why:



This time, the X’s just don’t line up well. They have no simple relation to each other. The fit just isn’t there. (Well, it is true that five time periods for C isn’t *so* far off from seven

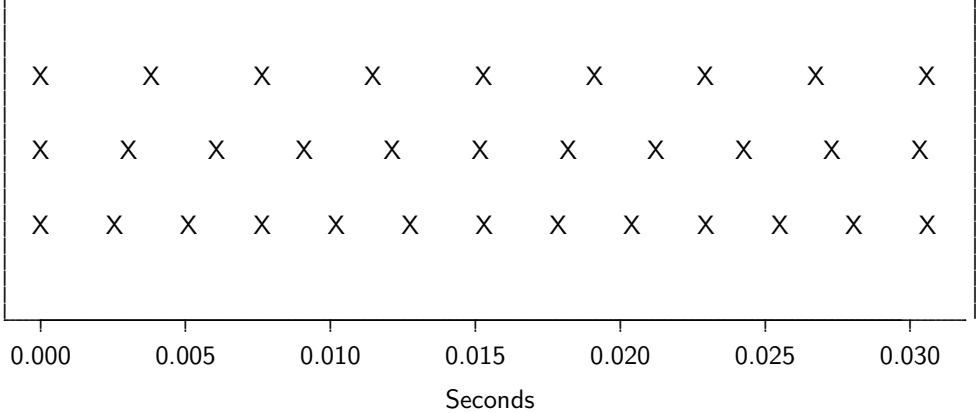
periods of F-Sharp. But it isn't so close either, and anyway, seven is too large a number of time periods to have to wait.) So, combining C with F-Sharp isn't so great, but combining C with G works well. These fundamental truths apply regardless of your favourite musical style.

**STRIKING A CHORD**

Combining the notes C and G produces a sound which is fine, but not very exciting. To get a really pleasing sound, let's add a third note – E. Middle E is four semi-tones (a “major third”) above Middle C, and three semi-tones (a “minor third”) below Middle G. When all three notes are played together, they form the “C major chord”, which is a sweetly harmonious, happy sound, like a barbershop quartet. It is the basis for music as diverse as *Row, Row, Row Your Boat* and the symphonies in C Major of Mozart and Beethoven and Schubert.

Why do these three notes – C, E, and G – sound so sweet together? Let's have a look.

**Arrivals for Middle C (top), Middle E (middle), and Middle G (bottom)**



We see that certain of the X's line up almost perfectly. Indeed, 4 time periods for Middle C, and 5 time periods for Middle E, and 6 time periods for Middle G, are all practically equal. So, about once every 0.015 seconds, the X's all line up.

So, even though certain of the X's do not line up for these three notes, quite a number of them do. This combination of variety and consistency is just what is required to produce the C Major chord – one of the most pleasing sounds known to humanity, and the basis of multitudes of tunes from Mendelssohn to Metallica.

So, humanity's quest for beautiful music amounts to finding creative and interesting ways to line up the air pocket X's of different musical notes.

## TEMPER, TEMPER

So why is it that the notes C and G fit together well, but the notes C and F-Sharp do not? Who decides which notes will have their X's line up well, and which notes will not? The answer is, mathematics decides!

Remember that High C has a frequency which is *twice* as large as Middle C. On the other hand, High C is twelve semi-tones above Middle C. There are eleven other notes – C-Sharp, D, D-Sharp, E, and so on – squeezed in between Middle C and High C. What are their frequencies? How do they fit in?

Early musicians – as far back as the Greek mathematician and musician Pythagorus of the sixth century B.C.E. – tuned musical instruments specially for each key. For example, when playing in the key of C, the note Middle G would be tuned so that its frequency was *exactly*  $3/2$  times that of Middle C. This would make the X's line up perfectly, so the notes would fit together exactly right. However, this system required different tunings depending on which notes you were planning to play, or in which key your music was written.

Over the last couple of hundred years, a more universal system has been used instead. This system – called *equal tempering*, a version of *well tempering* – spaces all twelve notes of the octave equally. In this way, a single tuning can be used regardless of the key being chosen, or the music being played.

Equal tempering is a system for breaking up each octave into twelve equal semi-tones. Since each octave represents multiplying the frequency by a factor of 2, each semi-tone represents multiplying the frequency by the *twelfth root of 2* – a number which produces two when twelve copies are all multiplied together. This number is about 1.059463, because twelve copies of 1.059463, all multiplied together, is equal to 2.

So what does this mean? If we start with Middle C as a base, then Middle C-Sharp, one semi-tone above Middle C, has a frequency which is 1.059463 times as large as Middle C. Then Middle D, two semi-tones above Middle C, has a frequency which is 1.059463 multiplied by 1.059463, or 1.122462, times as large.

Continuing in this way, we eventually reach Middle G, seven semi-tones above Middle C. The frequency of Middle G is 1.498307 times that of Middle C.

But wait. 1.498307 is almost exactly the same as 1.5, or  $3/2$ . This explains why every three time periods for Middle G, correspond to two time periods for Middle C. The two frequencies have a nice simple ratio,  $3/2$ , and that is why they fit together so well.

Similarly, the frequency of Middle E is 1.259921 times that of Middle C. This ratio is very close to 1.25, or  $5/4$ , and explains why every five time periods for Middle E correspond to four time periods for Middle C. Hence, like G, the note E also fits in well with C.

Another good example is the note Middle F, whose frequency is 1.334840 times that of Middle C. This is very close to 1.333333, or  $4/3$ . And, indeed, C and F also fit well together (their interval is called a “fourth”).

On the other hand, Middle F-Sharp has a frequency which is 1.414214 times that of Middle C. The number 1.414214 isn’t too close to any nice simple fraction (it isn’t *too* far away from 1.4, or  $7/5$ , but numbers like 7 are too large to be of much help). That is why the notes F-Sharp and C do not fit well together.

So, to figure out which notes fit well together, we don’t need to guess, or use trial and error, or study musical theory for years. We just need to remember the equal-tempering principle, and multiply copies of 1.059463 together to see whether or not we get a simple ratio of frequencies.

## UNLOCKING THE KEY

Experienced musicians are always discussing what *key* to play in. There is no cooler moment for a jazz musician than, when asked if they can play *I Got Rhythm*, to be able to reply, “Sure, dude. What key?”

Changing a song’s key raises or lowers all the notes by the same amount. This change might make the song more convenient to play on a particular musical instrument, or more comfortable for a vocalist to accompany. If done properly, the key change should have essentially no effect on the way the song “sounds” – it should be just as recognisable, just as lovely, and just as catchy, in the new key as in the old. Indeed, it should be exactly the same song, just performed at a higher or lower pitch.

How can this be? How can a song be the same, but different? How can we change the pitch without changing the tune?

From our understanding of frequencies, the answer is clear. To make the notes higher (say), we make their frequencies faster, i.e. make the X’s closer together. But to make the song sound the same, we leave the *relationships* between the X’s just as they were before. We compress each row of X’s by exactly the same factor. Equivalently, we raise each note by exactly the same number of semi-tones.


For example, suppose you are playing a song in the key of C, and are using a C Major chord, with the notes C, E, and G as before. You then want to change to the key of E, and use an E Major chord instead. How do you do this?

Well, to change Middle C to Middle E, you have to raise the note by four semi-tones, or equivalently multiply the frequency by 1.259921. So, to preserve the sound of the chord, you have to do *exactly the same thing* with the other two notes. Middle E also gets raised by four semi-tones, with frequency multiplied by 1.259921, to turn into Middle G-Sharp. And,



Middle G gets raised by four semi-tones to High B. The resulting air pocket arrivals look like this:

**Arrivals for Middle E (top), Middle G-Sharp (middle), and High B (bottom)**



We have thus created a chord E Major, which is just like C Major, except four semi-tones (or a factor of 1.259921) higher. The E Major chord sounds just as sweet, just as pleasing, as the C Major chord we saw earlier. The only difference is, the E Major chord has a somewhat higher pitch. We have successfully switched keys, from C to E.

There are many similarities between this E Major chord (consisting of E, G-Sharp, and B) and the previous C Major chord (consisting of C, E, and G). In each of them, every four spacings in the top row, and every five spacings in the middle row, and every six spacings in the bottom row, all line up well. So, every little while, the X's for all three rows are practically on top of each other. This is why the two chords have the same pleasing sound.

On the other hand, for E Major this confluence happens faster. While for C Major the X's line up about once every 0.015 seconds, for E Major they line up about once every 0.012 seconds. That is why E Major – while sounding similar to C Major in every other respect – also sounds somewhat *higher*, a fact the vocalist might well appreciate.

So the next time you are accompanying a beautiful starlette in her triumphant performance of her new jazz hit, and she turns to you and says with a wink, “Let’s do it in the key of E instead of C this time”, don’t panic. Simply raise all of your frequencies by a factor of 1.259921, and wait for the audience’s wild applause.