24th INTERNATIONAL CONGRESS ON SOUND AND VIBRATION
23–27 July 2017, London
ICSV24 23-27JULY 2017 LONDON CALLING

THE MIRACLE OF RESONANCE AND FORMANTS

Eric J. Heller
Harvard University, Department of Physics, Department of Chemistry, Cambridge, Massachusetts, USA
email: heller@physics.harvard.edu

Resonance underlies many phenomena in acoustics. Here we put resonance in its widest possible context, applying it to situations not commonly recognized as resonant phenomena. The reward is a greater understanding of acoustical function. In broadest terms, resonance results from adding amplitude constructively to pre-existing amplitude, be it sound, water waves, quantum waves, spring-mass oscillation amplitude, etc. Anti-resonance, involving destructive addition of amplitude, lurks nearby anywhere resonance rules. In fact there is a sum rule preventing net increase or decrease in power if the spectrum is averaged over a wide enough energy range. Resonance is the key to passive amplification. Speech formants are vocal tract resonances, not the survivors of filtering a source as they are so often portrayed. Resonance always finds itself in competition with friction, a battle that ultimately determines the timbre of instruments and the function of bells on wind instruments. The famous reproducing chamber of Edison phonographs is an example leading to a resonant "formant" that spans the whole audio spectrum of the device. Much the same can be said for a passive bullhorn mouthpiece. Most musical instruments use a "gatekeeper" (such as the cork upstream of the embouchure hole on a flute, or a trumpet mouthpiece) possessing a desired resonance or resonances early in the acoustic chain that imprint their broad peaks on the instrument's spectrum, shaping its timbre. Our hearing apparatus, including neurological components, is a superb detector of resonance, in the form of pitch and autocorrelation. Several examples are given, including repetition pitch. A simple and surprisingly accurate mathematical formula for perceived pitch in terms of strengths and frequencies of any number of partials is given, with unexpected side effects and consequences. Audio examples of these effects are provided

Keywords: (Resonance, psychoacoustics, pitch)

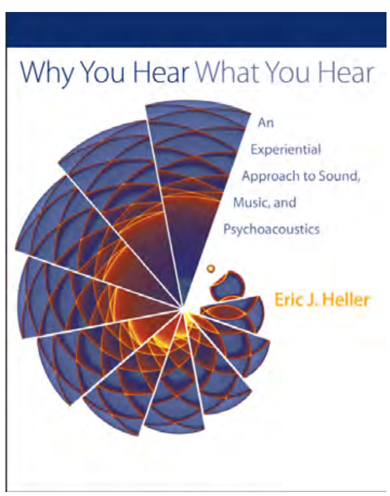

1. Resource

Some of the material contained in this article, and most of its points of view, are also found in a recent book written by the author, *Why You Hear What You Hear*[1], published by Princeton University Press. Its 600 pages and 400 figures do their best to explain acoustical phenomena (at a simple mathematical level) rather than merely describe them. The book is also peppered with new results, examples, audio files, and anecdotes, and associated with an extensive website with resources,

whyyouhearwhatyouhear.com.

2. The ubiquity and power of resonance

2.1 Amplitude and energy

Resonance happens whenever amplitude (of the kind that is squared to get power, probability, etc.) is added in-phase to pre-existing amplitude. The overworked canonical textbook example of the driven spring-and-mass oscillator is an example; there, the amplitude is velocity of the mass; in-phase forces add constructively to the velocity. The miracle of resonance stems from the fact that

$$\text{Energy} \propto \text{amplitude}^2, \text{ and}$$

$$\text{Power } \propto \text{amplitude} \cdot \frac{d\,(\text{amplitude})}{dt}$$

Thus doubling the amplitude a, i.e. $A = 2a = a + a$, gives four times the energy $4|a|^2$, but this is always counterbalanced by some nearby destructive anti-resonant addition, at worst nulling the amplitude, $A = 0 = a - a$. These two examples are the extremes, interpolated by the relative phase ϕ of the amplitudes, $A = a + e^{i\phi}a$. After averaging over all relative phases ϕ in the addition, doubling the amplitude just doubles the energy.

2.2 The reproducer chamber

Our first example of resonance involves the passive amplification achieved by the Edison reproducer, figure 1. The venerable reproducer on the Edison Diamond Disk gramophone has a flexible membrane driven by a thread in the middle; the thread is connected to the diamond tipped needle via an armature following the undulations in the recording medium. If the membrane instead lay across the throat of the horn, the sound would not be nearly as loud, nor of the same timbre. The chamber defined by the top hemisphere and the membrane is so narrow that pressure waves launched by the membrane reflect back and forth between it and the hemisphere many times during a typical one quarter oscillation of the membrane at audible frequencies. This automatically means constructive near instant addition of the amplitudes launched at different moments by acceleration of the membrane. Indeed the pressure build up in the chamber during a push inward by the membrane seems obvious, but the two timescales just mentioned need to be compared.

Engineers have an apt name for this: it's called loading the source. But under the wide tent of resonance described as constructive addition of amplitude, it is indeed and example of resonance. There is much more that could be said, probably much if it never investigated (after all, the reproducer is not a key piece of technology any more). There is the question of the size of the throat: how much of the pressure should be released? If the throat is too large, less constructive addition of amplitude, or pressure build up if you like, will happen, and the resonance will be weak. If too small, pressure build up is larger, but the amount of sound energy released per stroke versus that lost to various dissipative mechanisms begins to suffer. The best combination was probably found by trial and error (after all, it was Edison), but the result should be close to the following principle familiar from the theory of driven RLC circuits. The energy usefully extracted by the load will be maximal when it is equal to the intrinsic and unavoidable resistance in the device: wires, inductor, and unwanted resistance in the

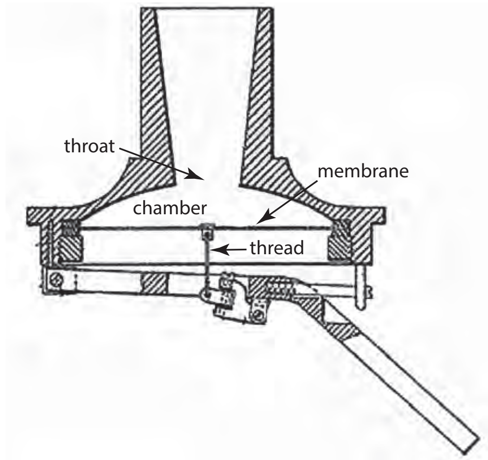

Figure 1: The Edison reproducer, from the patent application (the protrusion is the finger pickup),

load itself. This is the *maximum power transfer theorem*, or Jacobi's Law (1840): to obtain maximum external power from a source with a finite internal resistance, the resistance of the load must equal the resistance of the source.

This theorem does not seem to have the currency it deserves in acoustics. It affects strategies to make a singer's voice as loud as possible and the function of a trumpet bell. In both cases, the load is the sound energy being released, and the source is the vibrating vocal folds and pressure behind them, or the vibrating lips and pressure behind them. Friction follows from motion of air past walls, flexure of walls, and dissipation from net transfer of heat to the walls as air warms and cools adiabatically during the pressure changes during oscillation.

2.3 The Trouble with "Treble Boost"

The gradual enlargement of the trumpet tube after its final bend, culminating in a relatively sudden flair, is a highly refined shape with specific design goals for the resonance frequencies and timbre of the trumpet. The bell

- *assists the harmonic alignment of the instrument's resonances;*
- *controls the strength of the upper harmonics and thus the timbre of the trumpet;*
- *beams higher frequencies preferentially forward.*

The first effect of the bell is the most important: shift the higher harmonics, putting them closer together and into harmonic alignment, by effectively lengthening the trumpet as the frequency goes up. That sounds like a tall order, but the trumpet (or trombone, etc.) bell shape accomplishes this nicely, but reflecting hight frequencies further toward the end of the bell. Jacobi's theorem comes into play for the second of these functions.

The second effect of the bell is *supposedly* to enhance the strength of the upper harmonics by letting them out with little reflection (so-called "treble boost"), thus contributing to the brilliant timbre of the trumpet. Although it is noncontroversial that the bell allows higher frequencies to escape more efficiently, the net effect of this is subtle. Whether or not letting out the high frequencies actually makes them louder depends on several factors.

The third effect of the bell is not something we necessarily would have wanted: the higher frequencies are beamed preferentially forward, so that the timbre of a trumpet aimed at you and one aimed away is different.

The treble boost argument goes like this: the bell assists sound in leaving the trumpet by smoothing the impedance transition to free air, and giving an effectively larger tube to radiate the sound. A "cut-off frequency" for a tube, whether it has a flair or not, or toneholes or not, is the frequency above which the reflections are weak. A bell *lowers* the cut off frequency compared to a straight tube, i.e. makes it easier for high frequencies to leave by releasing them from an effectively larger diameter. The treble boost idea is that letting the higher frequencies out allows them to be better be heard, which seems simple enough.

The bell → treble boost → louder upper partials → brilliance of trumpet explanation sounds right and is easy to remember. But sound from a trumpet is not like beer from a spigot, where opening the spigot wider gets you more beer. The fatal flaw in the treble boost story is neglect of the effects of resonance.

Counter examples to the treble boost myth are all around. A french horn has a huge bell and yet is extremely mellow, even with the hand removed from the bell. There are only very weak upper partials. There is another problem with the treble boost idea: if it really worked as advertised, we wouldn't like it. The reason is psychophysical: we actually find that presence of a lot of very high partials in wind or most other instruments unpleasant. Playing a straight narrow tube using a trumpet mouthpiece indeed sounds raspy; high partials make it so. The distaste most of us have for strong high partials is presumably due to *self-dissonance*: adjacent partials falling within the same "critical bands" of our hearing. If a bell boosted the power throughout the treble range, we would not like the results. *The treble boost idea is backwards, what happens in mainly treble attenuation!*

The missing element in the overly simplistic treble boost idea, and the direct explanation of the two examples just mentioned, is resonance. If sound partially reflects from the open end rather than being fully released, the reflection returns to the lips to be resonantly reinforced if the lips are pulsing at the right frequency, making the sound louder. It can become so much louder that the release of a smaller percentage of it at the bell is more than compensated. Thus killing the resonance by letting the sound out on the first try *reduces* the power radiated. This is demonstrated with a numerical simulation in figure 2, where the sound emanating from a nearly pinched off tube is much louder than the same tube with a bell, even with the same drive. Care must be taken to adjust the frequency to the nearest resonance peak in both cases. This simulation had no wall friction included, and we discuss that shortly.

Sophisticated measurements of the treble boost using impedance and "pressure transfer function" concepts unfortunately had flawed logic. Investigating the trombone with and without a bell, the pressure transfer function $T = 20\log\left[p_0/p_i\right]$, where p_0 is the output pressure at the bell or end of the tube and p_i is the pressure in the mouthpiece was measured, reference [2]. There was an apparent treble boost in that measure, which not the same as loudness outside the bell. Nobody ever hears the power transfer function. They hear loudness of sound radiated by the bell.

Much more relevant is the average power delivered to the mouthpiece by a high impedance source. The average input pressure p_i cannot be used to calculate the average input power, and is insensitive to whether the pressure wave supplied by the source is in phase with the returning pressure from end reflections (if present) from the bell, but the power is quite sensitive to in-phase resonance, just as we have discussed. Moreover, if the power at the mouthpiece is doubled, it is clear the power output must be doubled, in the absence of significant nonlinear effects for any fixed set of partials. The pressure transfer function is a misleading measure, which can be proved by mundane measurement of loudness outside the instrument.

2.3.1 The battle between resonance and wall friction

Dissipation of sound energy at the walls (friction) is a crucial factor. If the instrument is fighting large dissipation it does help to let out more of the sound (less reflection) at the end of the instrument - because to reflect it back down the instrument is to lose more of it than you gain back through resonance. Thus resonance and wall dissipation are locked in a battle: they suggest opposite strategies

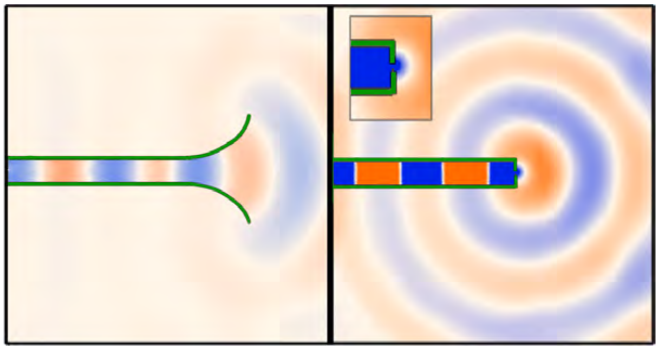

Figure 2: If the internal damping of the tube is low, a bell drastically *decreases* the power of a partial relative to a more reflective tube without a bell, or as seen here, more extremely, with a reflecting anti-bell. This is in spite of the much larger ratio of pressure inside to pressure outside in the case of a resonantly reflecting anti-bell. The insert shows the detail of the anti-bell. The color darkness represents wave amplitude; the total power, proportional to the square of the amplitude, is more than 10 times greater with the anti-bell in place. The forward projection of the wave by the bell is also apparent. The source was high impedance, adding fixed amplitude independent of what was returning by reflection.

for getting more sound out of the instrument. Jacobi's Law comes into play.

Jacobi's Law will adjudicate the competition of resonance, release, and dissipation. Assuming a 2.3 m round trip, sound at 2000 Hz, and an average radius a of 0.01 m, sound suffers about a 25% amplitude attenuation. This is greater than the loss due to sound radiation for all but very high frequencies if there is no bell. With a bell, equality of frictional losses and sound radiation losses a will move to lower frequency, because the bell allows escape of lower frequencies than before. The high partials find themselves losing sound intensity; *these partials are actually attenuated*; i.e. the bell thankfully *kills* the annoying higher partials that you hear in a cut off tube when it is played like a trumpet. They are diminished because of excessive sound radiation and failure to build intensity through the mechanism of resonance. This is the crucial detail omitted from the classic "treble boost" discourses. The near elimination of the resonances above 2000-2500 Hz when a bell is added is evident in figure 3.

Let's take a look at some hypothetical numbers. At 220 Hz, the wavelength is 80 times larger than a 2 cm wide tube. The sound is almost totally reflected. An 0.5 % end amplitude loss (sound escape loss) without a bell might become 1% with a bell, since the wavelength is still quite long compared to the bell diameter and the bell doesn't help much. This is a boost, but with 10% wall friction amplitude loss per round trip, only a small fraction of the sound at 220 Hz is getting out in either case. One the other hand at 1000 Hz, amplitude reflection loss (escape) of 2% with no bell might become 20% - freeing up much of the power which was lost to friction before the bell was in place. Finally at 8000 Hz, a 30% amplitude end reflection loss with no the bell is less than the 55% wall losses, but with a bell easily 97% of the sound could be radiated by the bell, i.e. amplitude reflection loss could plummet to 3%, since the smooth transition of the bell efficiently radiates high frequencies so well. This puts the 8000 Hz partial far to one side of Jacobi;s Law; power of a high resonant partial at 8000 Hz would be reduced by the presence of the bell, not enhanced by resonance.

2.3.2 Measuring the effect of adding a bell

Figure 3 compares sonograms for a plain tube buzzed with a trumpet mouthpiece, and the same tube with a funnel shaped bell added. The sound file, found on the CD for the excellent Springer

Acoustics Handbook, is intended to demonstrate treble boost upon addition of a bell. However once differences in overall loudness are accounted for, the bell has subtracted from the *relative* strength of the higher upper partials, as we have predicted. There is some evidence for mid-range boost. The same can be said for the set on the right, without and with a bell, in the author's analysis of Prof. Joe Wolfe's example files found on the informative University of New South Wales website.

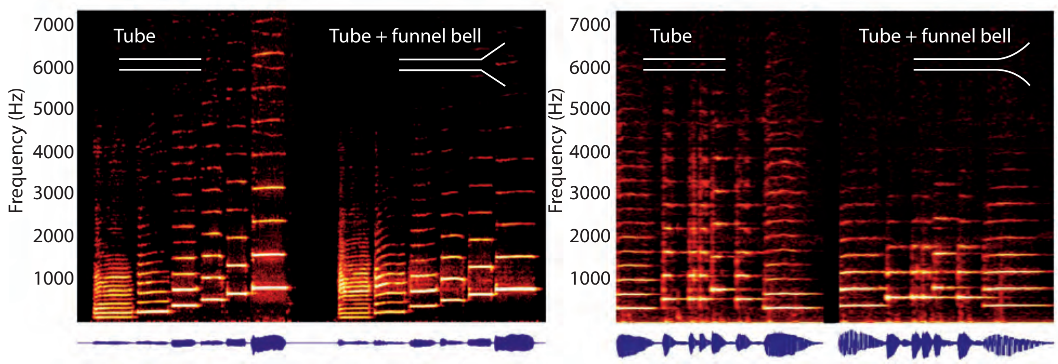

Figure 3: (Left pair) Comparison of partial strengths in six lip-buzzed notes before and after the addition of a funnel shaped "bell" to the end of a straight tube. The analysis was done on a sound file taken from the Springer Acoustics Handbook, meant to illustrate the increase of the intensity of upper partials due to a bell. After correcting the files to give the same loudness in dB, the bell is seen to have if anything *reduced* the strength of the higher partials relative to the lower ones, but some evidence of treble boost is seen around 800 Hz. (Right pair) Comparison of partial strengths in lip-buzzed notes (no mouthpiece) before and after the addition of a bell to the end of a straight tube; analysis of files found on http://www.phys.unsw.edu.au/jw/brassacoustics.html. Again, the bell has *reduced* the relative strength of the higher upper partials, but some evidence of treble boost is seen around 1000 Hz.

The treble boost idea is complicated, and needs more investigation!

2.4 Pulse puzzles

Consider two sequential but otherwise identical, non-overlapping sound pulses of duration τ, separated in time by $T > \tau$. The pulses are non-overlapping in both time and space; the total energy must therefore be twice the energy in one of them. There seems to be no chance for resonance here; the amplitudes cannot interfere in time or space. They can however interfere with each other in the frequency domain. This could be important to some object with its own natural frequencies receiving the pulses. The receiver could be us! We first consider the power spectrum of one of the pulses alone. The power it reveals integrating over all frequencies (area under the upper left curve) has to be half of what both will reveal together. So, where's the resonance?

The power spectrum (power plotted as a function of frequency) for a single Gaussian pulse is itself just a smooth Gaussian (the all-positive pressure profile in figure 4 can be achieved by smoothly accelerating and a decelerating piston toward a closed end of a tube. The resulting pressure wave will rebound from the end and return to the piston as the pulse shown. Likewise for two successive such accelerations toward the closed end before any return, resulting in the middle panel scenario when the pulses do come back.

There is a perceived pitch associated with the arrival of the two pulses, given by the inverse of the time delay between them. This is an example of a repetition pitch, the aptly named perception of a pitch when sound is repeated, even once, at audio frequencies. It is hardly a pure tone, but rather a vague sensation of a definite pitch that can be precisely matched with a variable pure tone selected freely by the observer.

On the other hand a pair of pulses of opposite sign, as on the right of figure 4, can be produced by accelerating the piston toward the closed end and later withdrawing it, accelerating in the opposite direction.. What repetition pitch, if any, will you hear? There is no positive revival of the autocorrelation at any time, only a negative one, and literally there is no tendency to repeat, unlike the case in the middle of figure 4. Will a pitch be perceived? To answer this, and emphasize the possibly

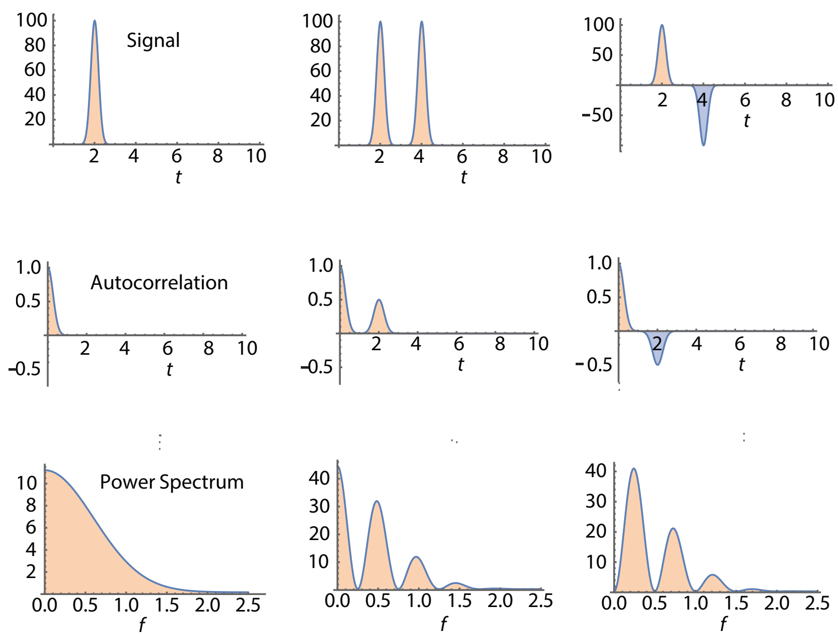

Figure 4: Examples of one pulse (left column), and two cases of two spatially and temporally separated sound pulses. In the left column, the pulse, its autocorrelation, and its power spectrum is shown. In the center column, the two pulses are individually identical in magnitude and duration with the one on the left, but they are not interfering with each other in space or time. Thus the total power in both must be twice that in one pulse. However, in the Fourier space, i.e. frequency, the two pulses interfere strongly. The inverse of their time separation (which is seen in the autocorrelation) is the frequency seen as the progression of peaks in the power spectrum in the center example. At the right, the peak spacing is the same but the fundamental is down and octave, with every other partial missing. Note that the total power (area under the curve) can be seen in both cases to be twice that of the single pulse power spectrum, if attention is paid to the vertical scale. All three cases are physically realizable in natural settings.

evolutionary component, suppose you are trapped in a large pipe. You are shuffling in one direction, and notice the noise you are making is forming a pitch. If the tube end is open, the sound you are producing is reflected back with an inverted waveform, delayed by the round trip time from the end of the pipe. If the pipe is closed, the sound is reflected just as it was produced.

The open end reflection is remarkably like the question we have posed: will you hear a repetition pitch from the inverted reflection? The answer is yes, but the pitch is lower by an octave than if the tube has a closed end. The sound will also posses a "hollow" timbre, so that even if the tube is dark, you would be able to tell it was open. The hollow sound is fascinating - indeed hollow half open tubes ought to sound hollow! The sound pulses emanating from the open end after an input pulse alternate in sign, and indeed as is well known the resulting spectrum is missing the second, fourth, etc. partials, like a the hollow sound of a clarinet.

3. A simple formula for pitch

It has long been known (but unfortunately not universally acknowledged) that the perception of a pitch or pitches in a sound is derived from its autocorrelation. The brain must be performing an autocorrelation, either in effect or literally. And why not? Listening to real things with real consequences for life or death for millions of years, what would we and our ancestors want to know quickly, as in a summary, about a sound? The three "executive summaries" of sound that our analytical brain serves up to our conscious mind are loudness, timbre, and pitch. Pitch is tendency of the source to repeat itself. Any such tendency is reflected as a peak in its autocorrelation at the time of the repetition tendency. Every digital sound analysis program that tracks pitch dynamically is in fact tracking one or more of several closely related autocorrelation measures. As an example, the pitch $\bar{f}$ of an orchestral chime (which had better be pretty definite for musical reasons) almost never coincides with one of the frequencies f_i of its various unevenly spaced vibrational modes that are responsible for generating individual sinusoidal partials. However the autocorrelation function $A(\tau)$ of the chime waveform $p(t) = \sum_i a_i \sin(2\pi f_i t + \delta_i)$, namely the average

$$A(\tau) =< p(t)p(t+\tau) >= \frac{1}{T}\int_0^T dt\, p(t)p(t+\tau) \sim \sum_i a_i^2 \cos(2\pi f_i t)$$

over some decent interval of time T will often have a dominant peak (dominant if it is the first positive and sufficiently strong peak at some time before $\bar{t} = 0.05$ seconds, but after $\bar{t} = 0.0002$ seconds. Then $\bar{f} = 1/\bar{t}$. The sense of more than one pitch present, as in a chord, can result is another strong or even stronger peak occurs after the first one.

Many works over the years have attempted to predict the pitch humans will sense, given a set of partials. A great many of them have considered a set of partial frequencies without concern for their amplitudes, which is clearly immediate grounds for dismissal of the proposed measure as as irrelevant.

We have found a simple, usually very accurate approximation for perceived pitch[1]. It seems to work very well in a reasonable range of circumstances. The idea is to approximate the time of early, tall autocorrelation peaks, given the set of amplitudes and frequencies as input. At the top of such an autocorrelation peak, the slope is zero. Therefore an approximation is developed for the time of such a zero. Its inverse gives the frequency.[1] Given a set of frequencies f_n and amplitudes a_n, (power $p_n = a_n^2$) the virtual pitch $\bar{f}$ which will be heard is given by

$$\bar{f} \approx \frac{\sum_n a_n^2 f_n^2}{\sum_n a_n^2 N_n f_n} = \frac{\sum_n p_n f_n^2}{\sum_n p_n N_n f_n} \tag{1}$$

where N_n is an integer depending on f_n: $N_n = [f_n/\bar{f}]$, where $[\ldots]$ is the integer nearest to the quantity inside the brackets; e.g. [5.53] = 6; [4.9]=5. This definition is slightly circular in that $\bar{f}$ depends on the integers N_n, which itself depends on $\bar{f}$, but in practice a self-consistent set of integers can be found.

As a test of the formula, we try the frequencies 820, 1020, and 1220 Hz of equal amplitude. The GCD of 820, 1020, and 1220 is 20, right at the threshold of hearing. This pitch seems an unlikely

[1] In calculus language, we take the derivative of the autocorrelation and set it equal to zero, $dc(\tau)/d\tau = 0$, and search near the recurrence we're looking for. Using the approximation $\sin(x) \approx x$, valid for small x, we get the formula **??**. This formula was first used in 1982 in the context of autocorrelations and molecular spectroscopy, by J. Zink, co-workers, and the author[4, 5].

perceptual result of combining these much higher frequencies. In his book *The Science of Musical Sound* John R. Pierce cites this case as an interesting example and reports that the perceived pitch 204 Hz; the formula 1 using $N_1 = 4, N_2 = 5, N_3 = 6$ gives 203.9 for amplitudes $a_1, \ldots = 1, 1, 1$. (Pierce did not report the amplitudes, but by experimenting with the formula it is found that the frequency is only mildly sensitive to them within reasonable ranges of equal amplitudes.) Figure 5 makes the situation clear. The autocorrelation function is shown as a thick black line, and the individual cosine terms contributing to the autocorrelation are shown in color. The small numerals near the peaks of the cosines counts the number of full oscillations starting at time equal zero. Near time $t = 0.0049$ the 820 Hz frequency has oscillated four times, the 1020 Hz frequency five, and the 1220 Hz frequency six times, thus $N_1 = 4, N_2 = 5, N_3 = 6$. A large peak rises at $t = 0.004904$, since all three cosines return to 1 near this time, although not exactly at the same time. The corresponding frequency is $f = 203.9 = 1/0.004904$ Hz. In spite of earlier recurrences, which would correspond to higher frequency pitches, this later recurrence is much stronger and dominates our sense of pitch. Precise measurement of the recurrence time from the autocorrelation function and formula **??** both give $f = 203.9$ Hz.

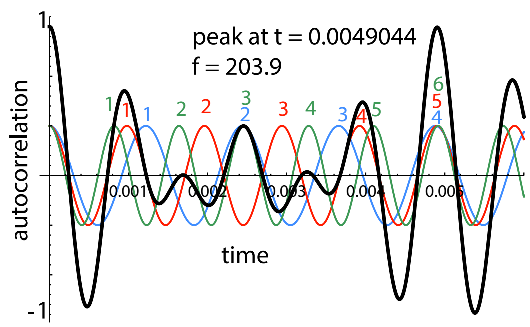

Figure 5: Autocorrelation function (black curve) analyzed for the perceived pitch corresponding to frequencies 820, 1020, and 1220 Hz with equal amplitudes. The autocorrelation function is the sum of the cosines shown in color.

A similar problem was considered by Plomp in 2001 (R. Plomp, *The Intelligent Ear*), using the frequencies 850, 1050, 1250, 1450, and 1650 Hz, which have a GCD of 50 Hz. Plomp reported that people perceive “about 210” Hz. The autocorrelation function peak suggests 209.2; and the formula 1 using $N_1 = 4, \ldots, N_5 = 8$ gives 209.13 for amplitudes $a_1, \ldots = 2, 2, 1, 1, 1$. Plomp did not seem to favor the autocorrelation idea; he advanced several other explanations for the apparent frequency shift.

We can check the autocorrelation formula against the class of examples suggested by Rausch and Plomp[3], who plotted the residue pitch (they called it the low pitch) against c, for the series $200 + c, 400 + c, 600 + c, 800 + c, 1000 + c$, for c on the interval $(350, 950)$. Their results, based on experiments with volunteer subjects, are shown in figure 6, taken from the article in Deutsch’s book, *The Psychology of Music*, along with our autocorrelation results. As c is increased, the appropriate integers N_i change, and are given at the top of the figure. The pitch obtained by autocorrelation (either by numerical peak finding, or from our simple formula 1- the results in this case differ by less than 1 Hz over the whole range) are shown in red. It is seen that the autocorrelation gives an essentially perfect estimate of the perceived pitch. At 1100, 1300, and 1500 Hz, there is an abrupt discontinuity in pitch, and at those frequencies the dominant pitch is indeed ambiguous.

The pitch formula equation 1 gives well known results, such as that the perceived collective pitch of two nearby, beating partials lies halfway between them (plus a small shift to higher frequency, which both the pitch formula and the actual autocorrelation predict). The “missing fundamental” or

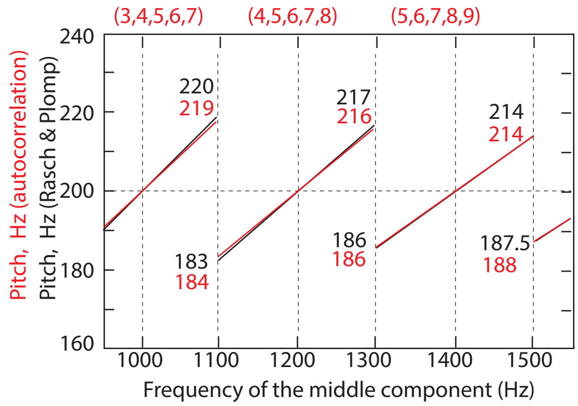

Figure 6: Residue pitch against c, for the series $200 + c, 400 + c, 600 + c, 800 + c, 1000 + c$, for c on the interval $(350, 950)$. Black lines and numbers: Results of Rausch and Plomp[3]. Red lines and numbers: autocorrelation results (both numerical and from formula 1 - they are very close to each other). The appropriate integers N_n to use in formula 1 are shown in red at the top.

"residue pitch" and myriads of generalizations for a harmonic tone with partials $f_n = nf$ are easily predicted, since

$$\bar{f} \approx \frac{f^2 \sum_n p_n n^2}{f \sum_n p_n N_n n} = f, \tag{2}$$

independent of the powers p_n, since $N_n = n$ in this case. One or more of the lowest frequency powers p_n can vanish without affecting the perceived pitch at the missing fundamental f.

3.0.1 A missing element

Perhaps the reader has noticed a conundrum: for time delayed and sign-reversed white noise added to itself, or two successive pulses of opposite sign as in figure 4 on the right, there is a negative peak in he autocorrelation after a time delay, but no positive one. Yet there is a pitch, that the pitch formula equation 1 is powerless to address without new rules. The hollow sounding pitch, as we have noted, is given by the inverse of twice the time to the negative peak. At this point we have no simple solution or extension to cover this case.

3.1 The Chorus Effect Blueshift

Two researchers independently found an interesting implication of the pitch formula, equation 1[6, 7]. Namely, that the perceived pitch of crowds singing the same note (or say the violins in an orchestra) with a range of off-key frequencies being sung or played (or perhaps due to independent vibrato) *lies slightly above the average pitch of all the sources*. This is evident from the formula, since the numerator goes as the square of the frequencies and the denominator is linear in the frequencies. For example, suppose suppose through error or deep vibrato the range of frequencies is clustered with a Gaussian distribution around the center frequency, with a 25 cent standard deviation. Then the perceived pitch will be 4 cents higher than the average pitch. Work on practical effects of this "blueshift" is in progress.

REFERENCES

1. Heller, E. *Why You Hear What You Hear*, Princeton University Press (2012).

2. Campbell, M., and Greated, C., *The Musician's Guide to Acoustics,* Oxford, pp. 347-350 (1987).
3. Deutsch, D., *Academic Press series in cognition and perception*, *The psychology of music*, (1999).
4. Tutt, L, Tannor, D., Schindler, J., Heller, E.J. and Zink, J. I., *Calculation of the missing mode effect frequencies from Raman intensities*, *J. Phys. Chem.* **87**, (1983).
5. Tutt, L.W., Zink, J. I. and Heller, E. J.,*Simplifying the MIME: a formula relating normal mode distortions and frequencies to the MIME frequency*, Inorganic Chemistry 26, 2158 (1987).
6. Baumgart, J., Personal communication.
7. Lacasa, L., *Emergence of collective intonation in the musical performance of crowds,* European Physics Letters 115, 68004 (2016).