Believing Your Ears: Probing the Brain through Musical Illusions

A CONVERSATION WITH

Diana Deutsch

Ideas Roadshow
INTELLIGENT. INQUISITIVE. INTERNATIONAL.

This eBook is based on a conversation between Howard Burton and Diana Deutsch that took place in La Jolla, California, on September 25th, 2014. This eBook reflects the unabridged conversation. The Introduction, Chapters 2a, 2b, 3a and 3b, are not included in the video version.

Edited by
Howard Burton

OPEN AGENDA
PUBLISHING

Table of Contents

Biography

Diana Deutsch is Professor of Psychology at the University of California, San Diego and a global authority on the psychology of music.

She obtained her B.A. in psychology, philosophy and physiology from Oxford (1959) and her PhD in psychology from the University of California, San Diego (1970), where she has held academic positions consistently since 1971.

A pioneer of harnessing computer-generated tones to carry out detailed aural experiments on music, memory, language and cognition, Diana uncovered a vast spectrum of musical illusions that are now standard in the literature, including the Octave Illusion, the Scale Illusion, the Chromatic Illusion, the Glissando

Illusion and many more. More recently, her research has focused on the role of perfect pitch in tone languages, for musicians and non-musicians alike, together with its implications for speech and music.

Diana has over 150 scientific publications, including *The Psychology of Music, Third Edition* (2013) and two CDs, *Musical Illusions* and *Paradoxes and Phantom Words and Other Curiosities*. Among her many awards and honors, she has received the Rudolf Arnheim Award for Outstanding Achievement in Psychology and the Arts and the Gustav Theodore Fechner Award. She is a Fellow of the American Association for the Advancement of Science, the Acoustical Society of America, the Audio Engineering Society, the Society of Experimental Psychologists, the American Psychological Society, and the American Psychological Association.

Howard spoke with Diana in her home in San Diego in September 2014.

Listening for Clues

In his celebrated *Essay Concerning Human Understanding*, the 17th century philosopher John Locke launched the modern empiricism movement in philosophy by enunciating the idea of a "tabula rasa" or "blank slate". All of our knowledge, Locke steadfastly maintained, comes to us through the act of sensory perception. For Locke, typically, this empirical perspective was not without its political and moral implications: any being free to choose his own formative perceptions was liberated from any pre-set "iron law of human nature" and was therefore free to "self-author its own character".

Nowadays, ethics, political theory and neuroscience are not quite as neatly overlapping, while modern science is much more likely to ascribe a significant role to innate, genetic factors for a wide range of human behavior.

Meanwhile, though perhaps no longer viewed as the sole factor in our intellectual development, sensory perception still naturally maintains its lofty perch as the key to understanding the world around us, with much contemporary cognitive science research devoted to clarifying just how our brains evaluate and interpret the information that is presented to them through our sensory organs.

Much progress has naturally occurred. After a tremendous amount of lively scientific debate, it is now generally accepted, for example, that the neural processing of vision occurs through two distinct pathways – the so-called "*what*" and "*where*" stream – that move through different regions of the brain.

Might something similar be happening when we hear?

Diana Deutsch thinks so. Diana is a renowned perceptual and cognitive psychologist at the University of California at San Diego, who has single-handedly discovered an astounding number of auditory illusions. Back in the 1970s, some twenty or so years before the "two-stream model of vision processing" was generally accepted, she came upon a very peculiar phenomenon known as the "Octave Illusion", revealing that different people would hear the same patterns of sound in strikingly different ways, largely depending on their handedness. After further experiments, she concluded that the likeliest explanation for this effect lay in a two-stream ("*what*" and "*where*") model of neural auditory processing.

A few people, like the famed Harvard neuroscientist Norman Geschwind, took notice, but for the most part Diana's findings fell on resolutely deaf ears. So she kept plugging away. After the Octave Illusion, came the Scale Illusion and the Chromatic Illusion. Then the Glissando Illusion, the Cambiata Illusion, the Tritone Paradox. The more she played with sound, equipped with state-of-the-art sound processing devices and an open, investigative perspective, the more she found. It turns out that tricking our minds to think it heard something it shouldn't have seemed frighteningly easier to do – a good deal easier, in fact, than creating comparative optical illusions.

Why was that, I wondered?

"I think that the auditory system is a messy one. If we think of it as beginning at the front end, there are something like about 125 million receptors in each eye and only 16,000 in each ear; and only about 3,500 to 4,000 of those in each ear are afferents, meaning that they actually send information up to the brain.

"So when you stack the numbers up, you can clearly see how dramatically low the number of receptors assigned to your ears is, especially given that the information they have to deal with is incredibly impoverished. Then if you consider the amount of brain tissue that's involved in vision as opposed to hearing, the difference is beyond enormous. In fact, if you think about it, it's really rather amazing that we're able to use our auditory systems at all to make much sense of the world."

I couldn't help thinking that this relative lack of auditory processing power represented an opportunity. If it was so easy to fool the brain into thinking that an objectively created sound was coming from somewhere else, or sounded like something else, then that seemed to open up plenty of occasions for probing just how our neurons processed auditory signals to begin with. Why weren't others rapidly hopping on the musical illusion bandwagon, anxious to see what else could be discovered?

"I think, partly, people tend to be theoretically-driven, and they also tend to be driven by what they expect, not just by the theory as written in the textbook, but what they really expect to hear. And when they hear something that's so absolutely crazy, such as the Octave Illusion, sometimes they don't like it.

"I have to say that I feel that myself too. There's a certain part of me that is made uncomfortable when I have to come to terms with some of these illusions. I don't really want to believe that the world is such a disorganized place. I'd like to feel that when I hear something, I know what I'm hearing and that other people are going to be hearing the same thing, and we can all understand it: one just needs to crunch through the various possibilities and it will all become clear. But when you come across something that's crazy, it's a bit upsetting. And I think that may be partly why not many people investigate auditory illusions.

"I think people find it difficult to articulate what they're hearing. So, if they hear a mush of stuff that doesn't sound quite right, they can just throw up their hands and blame themselves. In listening to the Octave Illusion, I knew that I was hearing a high tone intermittently in one ear and a low tone intermittently in the other ear; that's what I was really hearing and nobody was going to tell me otherwise."

Diana, you see, has the advantage of being born with perfect pitch, which has not only given her increased confidence in believing what she hears, but has also naturally driven her towards investigating the evolution and role of pitch in human speech, such as so-called "tone languages".

In yet another demonstration of how the plasticity of the brain enables it to adapt to external circumstances, her latest research demonstrates that those speaking tone languages have a much higher incidence of perfect pitch than others. Going even further, she speculates that there are interesting evolutionary links between music and speech.

"I do think that music and speech probably evolved from a proto-language that included absolute pitch, and that tone language is closer to that earlier language. Quite frankly, I think that non-tone languages such as English are in a sense degraded, in that they don't have tone involved, and I think that's a pity."

As John Locke maintained, then, it might well be that we all arrive in this world with a blank slate. But it turns out that we have a lot less choice of what to listen to – and how – than he might have supposed.

– Howard Burton

The Conversation

Introduction

Eclectic Beginnings

Music, art, philosophy, and a diversion into physiological psychology

Howard: I'd like to talk about your personal history and how you got involved in this line of musical research. Have you always been passionate about music?

Diana: Yes, music was a passion of mine from a very early age. As a child and teenager, I remember feeling sorry for people who were not immersed in music. I wanted to go to the Royal Academy of Music to study composition and piano, but my parents talked me out of it by saying that I'd never earn a living that way. It's a hard life; and even if I did earn a living, I'd have to travel all the time and would be very stressed and so forth. So they effectively talked me out of it, and I went to Oxford instead and studied psychology and philosophy, but I always felt that my heart was in music.

Howard: Piano was your primary instrument, I gather, but did you play other instruments as well?

Diana: While both my parents were very musical, neither of them played an instrument, and we didn't even have a piano around in the apartment that we lived in. However, when I was a very young girl, I had a little friend who lived in an apartment on a floor above us. She had a piano and I used to go over there and play around with it. At some point I must have been told the names of the notes, because I remember, when I was about four or so, being absolutely amazed that grown-ups had to go and stare at a keyboard in order to figure out what note was being played when they heard it.

I couldn't understand how people couldn't just hear a note and know exactly what it was right away. I remember being amazed, going from one person to another, to prove that they couldn't do it. At any rate, I wouldn't have had any musical training at that time had it not been for this little friend of mine. I started serious musical training when I was six or so, and immediately just plunged into it. That was when I really fell in love with music, and started practicing many hours a day.

Howard: And your parents were supportive of this?

Diana: They were, indeed.

Howard: Where were you, exactly? Which part of England did you live in?

Diana: I was in London at the time, and then I went to boarding school. It was a very strict boarding school with several negative qualities, but I was fortunate that they were also into music. As a result, they often let me out of class so I could go into the music rooms and practice the piano there. I did have good teachers there, and for that I thank them.

Howard: Where was the school?

Diana: The school was called Christ's Hospital or The Bluecoat School. The girls' school was in Hertford and the boy's school was in Horsham. Perhaps you've seen pictures of a Bluecoat boy dressed in a long, blue gown with yellow stockings[1] and so on? The girls didn't have that kind of costume, but we wore tunics and all that. The boys and girls met once every hundred years—

Howard: Once *every hundred years*? How very English.

Diana: Very much so. The school was founded in the 16th century; and indeed, we met once every hundred years. I remember we were given actual dresses to wear, and stockings, and we all trooped out to Horsham to meet with the boys, which was a very big deal.

Howard: I can well imagine.

Diana: Yes. And after five years there, I felt I couldn't take it any longer and I quit. I just went to a local school in London where I pretty much did my own thing. I was lucky to have a very sympathetic principal who let me not show up if I didn't want to; and as a result I spent an enormous amount of time on music.

Howard: Growing up in London, I can imagine that you would have had all sorts of opportunities to listen to music, to go to concerts and to be engaged in a broader musical life. Were you able to partake in all that?

Diana: Yes, to some extent. But I was also heavily involved with visual art: my father was a sculptor, and I spent a lot of time at the British Museum studying art there, as well as The Tate Gallery and The National Gallery. Also, I spent a lot of time watching my father at work. More than anything else, really, I think that my strong feeling that science and music ought to be drawn together sprung from my discussions with him and from his views. He was

1 See here or here, for example.

very, very thoughtful; and I found his work very meaningful. So that was a tremendous influence too – visual art. I think one can see that in my work, especially in the work that I did on musical illusions, because a lot of them have analogs in visual art.

Howard: Despite the fact that your father was a sculptor, he still deterred you from a career in composition and music.

Diana: Well, yes; because it was a very hard life and he knew that from personal experience. He was determined to be an artist pure and simple: he would not do commercial art – and, as a result, he never had any money. So he felt that if I went into music, then I would never have any money either. He was probably right, too. That was their only concern, really, the economics of it.

Howard: Was it difficult for you to then suddenly back away from music and go to Oxford to study philosophy and psychology?

Diana: Well, I was always interested in philosophy and psychology anyway – primarily philosophy, which I had studied alone quite intensively as a teenager. However, in those days at Oxford, you had to study philosophy along with something else, so I started off with philosophy, politics and economics (PPE) and then switched to philosophy, psychology and physiology (PPP), which was a very new program at the time.

Howard: When you spoke earlier about how you had been captivated by philosophy as a teenager, what sorts of things are you referring to? What had particularly stimulated you?

Diana: Well, at the time, I suppose big questions in philosophy like, *How do we know that we exist?* or *How do I know that other people have thoughts?* All those types of questions. Of course, they continued to intrigue me as I progressed through my studies at Oxford.

Howard: Did you continue to play music throughout your Oxford years? Did music still play a significant role in your life?

Diana: Yes, indeed. I had a piano at home, and I did a lot of playing and improvising on the side. But at the time, I didn't see it as feasible to be studying music formally in conjunction with psychology and philosophy; it just didn't seem possible, and I accepted the fact that music had to be a hobby.

However, when I arrived here at UC San Diego, people were just beginning to write software for generating music – in particular, Max Mathews[2] at Bell Labs, wrote terrific software for the generation of music. But I was able to use a different method. I used a computer to control a function generator so as to produce sine-wave tones in sequence, and the sine-wave tones were very well-specified in frequency, amplitude, duration and pause. In that sense I was very lucky, because I began doing this type of work at a time when it was only just becoming possible.

Howard: I'd very much like to get back to that, but I'm going to stubbornly return to Oxford, because I'm still imagining you as an undergraduate, and all of a sudden you're at UCSD, so I need to fill in a few gaps here.

So, you're playing music as a hobby while at Oxford studying psychology, philosophy and physiology – which seems like a rather broad spectrum, by the way –

Diana: Well, we didn't actually do all three of the "P's", as it were. You had to choose two of them.

Howard: Which one did you leave out?

Diana: I left out physiology, but I sort of studied it anyway. In fact,

2 Max Mathews (1926-2011), American computer music pioneer. See chapter references for details.

after I graduated, I wrote a book, together with my husband[3], called *Physiological Psychology*, which was the first textbook of its sort at the time.

Howard: That hardly sounds like leaving it out, then.

Diana: No, not really. However, it wasn't part of my formal curriculum. Just to fill in more of the gaps for you, I did, in fact, go to Stanford initially, after I graduated from Oxford, to take graduate courses in philosophy.

However, at the time, I really wanted to have a family and wasn't so interested in an academic career. So instead, I wrote this book together with my husband. We also wrote an article together on attention[4], which is still widely quoted, invoking something called "the late-selection view of attention".

Howard: How did all that happen?

Diana: Well, the *Physiological Psychology* part was really due to the fact that my husband was offered a contract to write this textbook, so I said I'd write it with him. It wasn't actually a particular interest of mine, it just happened to be convenient.

Howard: Okay, but is it fair to say that, through that experience, you started becoming interested in attention?

Diana: Well, that was a separate thing really, this whole issue of how attention works and questions about whether or not it *does* actually work. It had been believed at the time that there was a low-level filter that blocked out all except for – in the case of, say, sound – certain items in a certain frequency range, or items coming from a particular location in space – a very crude type of fil-

3 J. Anthony Deutsch, Professor of Psychology Emeritus, UCSD.

4 "Attention: some theoretical considerations" *Psychological Review* **70** 80-90, 1963.

tering. What we argued was that that wasn't what happened, but rather all the information went up to a very high level (indeed, the highest level) where it was analyzed for semantic content and so on; and at that point some high-level decision was made as to what to attend to. The information that was then deemed "important" was kept, while the rest was actually inhibited from consciousness. That is, in essence, the "late selection model".

Howard: I see. And where in the brain is that processing taking place, exactly?

Diana: We assumed that it was at the highest level in the cortex, as opposed to a very low level.

Howard: Have there been studies since then?

Diana: Yes. Since then there have been many articles on the late-selection versus the early-selection models of attention, using all kinds of paradigms, examining matters from a physiological point of view and so on, but I don't think there can be very much argument any more for the early-selection model.

At any rate, that's what I was doing for some time, together with taking some graduate philosophy courses at Stanford. And then my husband got a professorship at UCSD and I came with him and went back to graduate school there, this time in psychology. My heart was not really in doing physiological psychology, which would have been the obvious thing to do, given that I had just collaborated on writing a book on the subject.

Further References

Max Mathews (1926-2011) was a computer music pioneer who worked at Bell Labs for many years, after having graduated from MIT with a doctorate in electrical engineering. He is perhaps most famous for writing MUSIC (and many of its many subsequent descendants), a sound generation program, in 1957. For more background see *Electronic and Computer Music, Fourth Edition* by Peter Manning (2013), *Electronic Music* by Nick Collins, Margaret Schedel and Scott Wilson (2013) or *Introduction to Computer Music* by Nick Collins (2009).

Current perspectives on psychological attention models can be found in, for example, *Fundamentals of Cognitive Psychology, Third Edition* by Ronald Kellogg (2015) or *Cognitive Psychology: Connecting Mind, Research and Everyday Experience, Fourth Edition* by E. Bruce Goldstein (2014). See also Daniel Goleman's book, *Focus: The Hidden Driver of Excellence* (2013).

Chapter 1

Tones, Pitches and Critical Values

Intriguing results in music and memory

Diana: Soon after I arrived at UCSD, some people wrote software for generating sequences of tones that could be very carefully specified in terms of various parameters. So it then became possible, for the first time, to ask questions such as, *How does memory for pitch work? How does memory for duration work? How does memory for loudness work?* and so on, in a much broader setting than had been possible before.

Psychoacoustics, at the time, consisted mostly of answering questions like, *What is the softest sound that one is able to hear at a particular frequency?* or W*hat is the shortest sound that one is able to hear?* or *What happens when you play more than one tone at a time?*

However, all those kinds of questions weren't really relevant to music, as you couldn't find out much about music with these very simple sounds. Now, suddenly, the field was drastically opened up by the use of computers that enabled people to generate precise streams of sound. The problem with trying to study music without such precision is that your results are difficult to interpret.

Howard: There's no way to analyze it analytically or clearly.

Diana: That's right – or to demonstrate that you've really found what you think you've found, because there are just too many alternative possibilities. So I grabbed the opportunity provided by the new technology and started to study memory for the pitch of a tone in a sequential setting.

This is how these experiments would work. There would be a test tone, a sequence of other tones, and another test tone; and the subject would be told to forget about the intervening tones (that they're "just there"), and to judge whether the test tones were the same or different in pitch. The test tones would either be identical in pitch or they would differ by a semitone, which is the difference between one tone and the next one up or down on the musical scale (from C to C#, for example).

And I found, to my great surprise, that this was a very difficult task. It turned out that most people were close to chance in doing it. As a matter of fact, I have a couple of examples posted, and I think that it would be good to play those.

Howard: Sure. So, let me understand: the point is to try to remember what the first pitch is and compare it to the last one?

Diana: Yes. You're going to hear a tone, followed by a sequence of six intervening tones, which you should ignore – just try to remember the first tone. Then, there will be a two-second pause, followed by another tone that is either the same in pitch as the first tone or a semitone removed, and you have to decide whether the first and last tones are the same or different in pitch. You might expect this to be very easy, especially since you're told to ignore the intervening tones. (I mean, suppose you had two tones separated by five seconds of silence and you had to decide whether they're the same or different, that task should be easy for anyone to do.) The intervening tones, I should mention, are different in pitch from either of the test tones.

Diana: I would guess that, whether you got them right or not, you certainly would not have found them trivial.

Howard: No, indeed I did not.

Diana: Right. Again, this was a big surprise, because people would have expected that if they were told, “Don’t attend to the intervening tones,” then these tones wouldn’t have any effect on their judgments of the test tones. But in point of fact, most people find it extremely hard to ignore the intervening tones.

If there are eight intervening tones rather than six, you’re really close to chance, so with six I made it just a little bit easier. But then things got really interesting. Suppose, instead of an intervening series of tones, you have an intervening series of spoken numbers. You are told not only to say whether the test tones are the same or different in pitch, but you also have to repeat all the numbers that were played. It turns out that most people have no problem judging the test tones as the same or different, and at the same time recalling the numbers.

Howard: That was certainly much easier, as you say. One has a sense that there’s a certain interference that’s happening with the music somehow. You’d like to think that you’re able to remember something and put it in some storage space that is incorruptible, but it feels like the extra musical information that you’re getting in the first example seems to somehow corrupt our musical memories.

Diana: Yes. This obviously wouldn’t happen if the test items were numbers. There’s just something about memory for tones that’s causing this to happen. So I concluded that tones must be stored in a specialized store in which they interact with each other.

Then I went ahead and did a number of other experiments, where I showed that, depending upon the pitch relationships between the tones in the intervening sequence and the test tones, the error rate would go up or down. In this case, subjects were chosen who had very close to 100% accuracy in judging whether two tones were the same or different in pitch if there were no tones inserted during the time interval between them.

Fig. 1 - Percentage of errors in pitch recognition, plotted as a function of the pitch separation between a critical intervening tone and the first test tone. The peak of errors occurred when the critical tone was 2/3 tone removed from the first test tone.

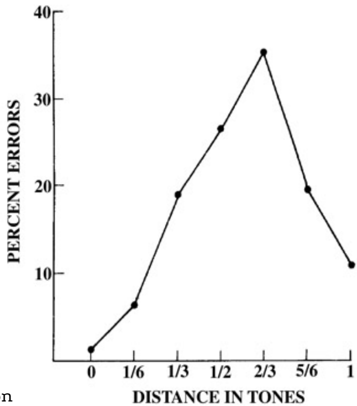

And it turned out that, if a particular interpolated tone was presented in the sequence, the error rate was a very clear function of its relationship to the first test tone. For example, when it was two-thirds of a tone removed from the first test tone, error rates were very high, just due to this one intervening tone.

Howard: Okay, hold on for a moment and let me see if I understand this. The idea is that you hear one tone, then you hear a bunch of other tones, and then you compare that first tone to the last one; and if any of those intervening tones are sufficiently close to the original tone, then the likelihood that one would make a mistake in the experiment increases. There's a dependency.

Diana: Yes, there's a dependency; and depending upon what the actual relationship is, the error rate will be higher or lower. So, for example, (referring to Figure 1) when the critical intervening tone was only one-sixth of a tone removed from the first test tone, the error rate was very low; and when it was five-sixths of a tone removed from the first test tone, the error rate was also much lower than when it was two-thirds of a tone removed. And, it turns out that this must be due to an effect known as "lateral inhibition."

This had been shown in vision, whereby, for example, if you present a flash of light together with another flash of light, neurons that respond specifically to the first flash will fire either much more or much less, depending upon the distance between the two flashes. Lateral inhibition occurs because neurons responding to the second flash affect the neurons that respond to the first flash.

Our experiments showed that when the critical intervening tone is two-thirds of a tone removed from the first tone, which is a little over a semitone, the error rate is the highest. We confirmed this in several ways.

For example, we found a group of people who had extremely low error rates to begin with, and plotted their error rates in making these pitch judgments as a function of the pitch relationship between the first test tone and a critical intervening tone. We then did another experiment where we plotted their errors as a function of yet another tone whose relationship to the first, critical intervening tone also varied. And we found that when this second critical intervening tone was placed two-thirds of a tone from the first critical intervening tone, you can inhibit the inhibiting tone, and thus cause the error rate to go back down again.

Howard: That naturally makes me wonder what's so magical about this "two-thirds of a tone" differential between tones.

Diana: Well, from the distances in pitch involved it could relate to something called the "critical band," though this still remains to be determined.

Why is this important? Well, lateral inhibition is important in sharpening an image because you don't want to have factors in play that are going to cause blurriness. And I'm assuming that the same thing is going on in memory for pitch. So, in this case, you have a sharpening of the memory for the first test tone as a result.

Howard: I'm thinking that there are people who rely upon differences in tones more than others. For example, there are people who use languages that depend on tonal structure and pitch and so forth. I'm hardly an expert on this, but my understanding is that some Asian languages, for example, depend on this explicitly.

So, if we took, say, someone who spoke some language that depended explicitly on tonal structure, I would imagine that two things would happen. Firstly, I would imagine that you might find that subjects in this category who would be participating in these sorts of experiments might be better able to distinguish tones through experience, and I would also imagine that the language itself would carry some aspects of this structure with these particular ratios that you're talking about.

Diana: That's a very interesting observation. I think that there might be an effect of musical experience, and also an effect of language. Concerning your question about tone language speakers and how well they would do on memory for pitch, it just so happens that we're doing that very experiment right now.

Howard: Oh, really? So what have you found out? Do you have any data yet?

Diana: Well, the few tone-language speakers we have tested have done very well so far, but there have been too few to say for sure. But I wouldn't be surprised if they continued to do very well, because they are, indeed, very pitch-oriented.

Howard: Very interesting. But returning to the main point you raised earlier: I obviously don't pretend to know anything about this, but this business of a two-thirds tone difference strikes me as a really shocking result. I would expect something to go smoothly in one direction or another, but the idea of some mathematically precise critical point is quite counter-intuitive. That must have been really exciting for you.

Diana: Yes, indeed! Most of my results have come as a surprise to me. I certainly didn't expect this, but I just noticed it. It certainly was not predicted from any theory at that time, and it does show that memory for pitch is the function of a very specific store, or "module", and that specific interactions take place within this store, rather than memory for pitch just "being in there" along with memory for other types of material.

Further References

For more on the phenomenon of lateral inhibition see, for example, *Sensation and Perception* by Steven Yantis (2013) or *Sensory Neural Networks: Lateral Inhibition* by Bahram Nabet and Robert Pinter (1991).

The notion of a "critical band" was introduced by the American physicist Harvey Fletcher (1884-1981), who was involved in the famous "oil drop" experiment to determine the charge on an electron that helped garner Robert Millikan the 1923 Nobel Prize in physics (see his 1982 *Physics Today* article "My work with Millikan on the oil-drop experiment"). For more background on critical bands, see, for example, *Signals, Sound, and Sensation* by William Hartmann (1998) or "Critical Bands" by Bertram Schaf, chapter 5 of *Foundations of Modern Auditory Theory, Vol. 1* (p. 157, 159-202), edited by Jerry Tobias (1970).

Those of a more technical persuasion might be interested in the articles "Complex tone processing in primary audio cortex of the awake monkey. II. Pitch versus critical band representation" by Fishman et al. (*J Acoust Soc Am*. 2000 Jul;108(1);247-62), "Laminar fine structure of frequency organization in auditory midbrain" by Christoph Schreiner & Gerald Langner (*Nature* **388**, 383-386, 1997) or "Psychophysical tuning curve and critical band determined by masking in the presence of FM sounds" by Isojima et al. (*Auris Nasus Larynx* 1995;22(1):16-23).

Chapter 2

The Octave Illusion

How to easily confuse the brain with tones

Howard: What effects did these particular results have on other colleagues?

Diana: Well, they were definitely taken aback. That first experiment did get published in *Science*[5], so obviously the reviewers liked it. Many people working on memory were very surprised and said, "It can't be true." But then, of course, they tried it and it was true, so, eventually, people stopped arguing about it.

Howard: But I'm also wondering about its impact in terms of opening up a new field. This is a result that is not only counter-intuitive and exciting, but suggestive of other potentially unexpected results. You mentioned how you were applying this new technology of precise mathematical modeling of tones that hadn't been done before because the technology simply didn't exist, and I can imagine other people saying, "Well, what other effects might we have? How might we move forwards in this entirely new field?"

Diana: I'm trying to remember, but I don't think that, at the time, other people were using this particular technology. As I was saying, we controlled function generators to produce the sequences of tones. (These were made by Wavetek[6] in San Diego, which was very useful). Meanwhile, other people who were beginning to study sequences of sounds used the method of splicing pieces of tape together, but that was very difficult and time-consuming, whereas I could just type into the computer what I wanted in terms of frequency, amplitude, duration and pause for any num-

5 "Tones and Numbers: Specificity of Interference in Immediate Memory" (*Science*, 1970, 168(3939), 1604-1605).

6 Wavetek became Willitek and was later purchased by Aeroflex.

ber of tones. I could go on for hours, and the function generator and computer software would do the rest. But at that particular time, I don't remember anybody else using that technology.

Howard: Well, that's odd; but it's true that new technologies are sometimes overlooked for a while when they first come on the scene. At any rate, you found this intriguing result and were presumably keen to look for others. So, how did you move forwards?

Diana: At one point, I thought, *Well, I've got one stream of tones, maybe I could try a second stream of tones using two function generators.* I had a very good computer programmer who created software that would enable me to present two streams of tones at a time. Because of my earlier work on attention I was very interested in a body of literature where you would present one stream of speech to your left ear and another to your right ear, and people found results such as, if you were right-handed, you remembered more of the speech that came to your right ear than your left ear – that kind of thing.

So, I figured that I could try that same sort of experiment with musical tones. But as soon as I set this up, I came across an illusion that I found very difficult to believe. I presented a sequence of tones – a high tone and a low tone – that continuously repeated in both ears simultaneously, except that when the left ear got the high tone, the right ear got the low tone, and vice versa.

I was interested in finding out what the result would be. Would we hear the high tones going back and forth from ear to ear? Would we hear the tones coming to the right ear alone or to the left ear alone? What would happen?

When I put the earphones on and listened, I was absolutely amazed, because I didn't hear either of those things. Instead, the right ear seemed to be receiving an intermittent high tone while, at the same time, in alternation, the left ear seemed to be receiving an intermittent low tone. This wasn't at all what I had typed into the computer. What I had typed in was a continuous sequence of high tones alternating with low tones in one channel and another continuous sequence of high tones alternating with low tones in the other channel. But what I was instead *hearing* was "*high tone – silence – high tone – silence*" in one ear and "*silence – low tone – silence – low tone*" in the other ear. My immediate reaction was to check the computer; and when I found no problem there, I tried switching the headphones around. I found, to my amazement, that I continued to hear the high tone in the right ear and the low tone in the left ear.

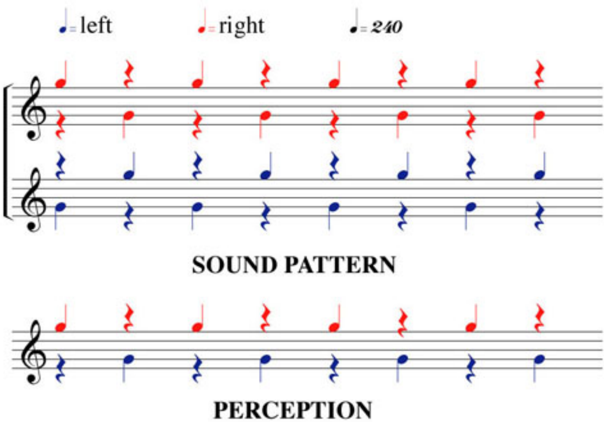

Fig. 2 - The pattern that produces Deutsch's Octave Illusion, and a way that it is often perceived when played through stereo headphones. The listener hears a high tone in the right ear that alternates with a low tone in the left ear.

So then I realized that this percept must be the product of my perceptual system, which was quite astounding. Today, people have heard of the octave illusion and you see it in textbooks, and so on. So when I try to explain how surprised I was to a group of students, they're not nearly as amazed as I was back then. Nevertheless at the time, it just seemed simply unbelievable.

Howard: Well, you must have then gone to somebody else and said, "Try this."

Diana: That's exactly what I did. I went out into the corridor and I yanked people in and played the pattern to them. I asked them what they heard, and they typically said, "That's rather strange," or just laughed or something.

By the end of the day, I'd ended up testing probably dozens of subjects. I noticed that people who were left-handed didn't necessarily hear the same thing as the right-handers did. Some of them would hear the high tone on the left and the low tone on the right consistently, and for others, the high and low tones would appear to switch ears, while yet others would hear rather complicated things such as a low tone that just seemed to move up and down in semitone steps in one ear and an intermittent high tone in the other ear.

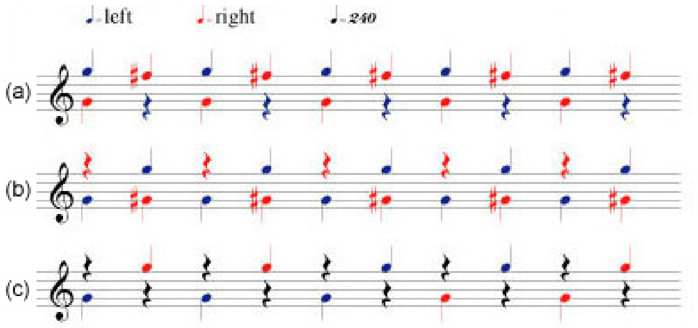

Fig. 3 - Some other ways that Deutsch's Octave Illusion is perceived. Some listeners hear a high tone that alternates from ear to ear with a slight change in pitch, together with an intermittent low tone in one ear, as in (a). Other listeners hear a low tone that alternates from ear to ear, together with an intermittent high tone in one ear, as

in (b). Yet other listeners hear a high tone in one ear that alternates with a low tone in the other ear, but the high and low tones often appear to exchange locations, as in (c).

Some people would hear the high and low tones alternating, but would not be able to tell which ear was hearing what, rather like a Necker cube in vision, or those Escher[7] lithographs, where you can suddenly change your percept.

A few people didn't hear a pitch difference between the tones, but I later came to realize that many of these people were not familiar with the octave relationship, and so did not recognize it and instead thought that they were hearing the same tone all the time.

There were several conclusions. First of all, the basic illusion that most people heard was very bizarre. Secondly, the correlation with handedness had to be important, because it had to show something about cerebral dominance. And thirdly, it opened up this whole question of, *When we're listening to musical patterns, are different people hearing different things?*

The simplest way I can describe what is going on is to say that you hear the tone that is coming into the dominant ear – which is generally the ear that matches your handedness – but you localize the tone to the ear that receives the higher frequency regardless of whether a higher or lower frequency is, in fact, perceived.

Howard: I'm not sure I quite understand that. Is it like one is establishing a bass line somehow?

Diana: Well, if you play both tones at the same time and you mix the channels together, you obviously hear just a single, low tone because, since they're at the octave; it's going to sound like the low tone together with the first upper harmonic.

7 M.C. Escher (1898-1972), Dutch graphic artist. See chapter references for details.

Now, a person who hears the pattern of tones that's presented to her right ear is going to hear the same pattern as this:

[play one channel]

and the person who hears the pattern of tones that's presented to the left ear is going to hear this

[play the channel again, but leave out the first and last tone].

But both sets of people localize each tone to the ear that receives the higher frequency, regardless of whether they're actually hearing the higher or the lower frequency.

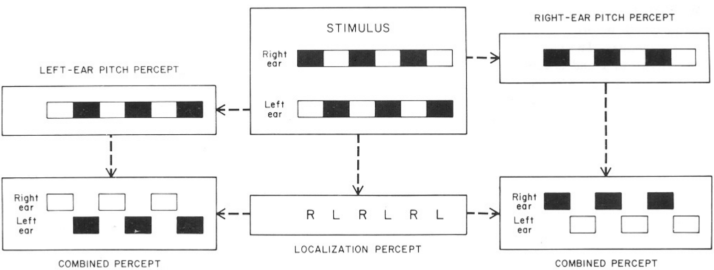

Fig. 4 - Model showing how the output of two decision mechanisms, one determining perceived pitch, and the other determining perceived location, can combine to produce the Octave Illusion. The black boxes indicate high tones and the white boxes indicate low tones.

For example, for a right-eared listener, when the right ear gets the high tone, and the left ear gets the low tone, they'd hear the high tone since this is presented to the right. On the other hand, each tone is localized at the ear that receives the higher frequency, regardless of whether the higher or the lower frequency is in fact perceived. But when the right ear gets the low tone and the left ear gets the high tone, the listener hears a low tone, since this is presented to his right, but he localizes it in his left ear, since this ear is receiving the higher frequency, So when you put these two decision mechanisms together (that is, the mechanism determining whether a high or low tone is heard, and the mechanism determining where the tone appears to be located), and play this continuous pattern, the listener hears a high tone on the right alternating with a low tone on the left. This is the typical Octave Illusion for the right-hander. On the other hand, because the localization rule is constant, a person who is left-eared is going to hear the high tone on the left alternating with the low on the right.

One of the things that the octave illusion shows is that there are separate pathways – one determining what pitch we hear (the '*what*' pathway) and the other determining where the sound is located (the '*where*' pathway) – and at some point their outputs come together. In this case, the combined action of the '*what*' and '*where*' decision mechanisms produces the Octave Illusion.

So to repeat, when the right-eared listener hears a *low* tone in the *left* ear, he's actually *receiving* a *low* tone in the *right* ear and a *high* tone in the *left* ear. This is then a case of where the auditory system goes wildly wrong, hopelessly wrong, with a very simple percept.

The bottom line is that the auditory system is really a very messy system that fudges signals through it, to the point that we do usually hear things correctly – but in order to do so, we need a lot of top-down processing to clean up the signal and make sense of it. When the auditory system is prevented from making corrections in a simple experiment like this one, it can make absurd errors.

Further References

More background on the octave illusion, with references and further analysis, can be found on Diana's Octave Illusion page of Diana's general website.

M.C. Escher (1898-1972) was an influential Dutch graphic artist. For more information about Escher see this website published by the M.C. Escher Foundation and the M.C. Escher Company B.V.

Studies of Escher and his work include *M.C. Escher: Visions of Symmetry* by Doris Schattschneider (2004) and *M.C. Escher: His Life and Complete Graphic Work* by Bool et al. (1992). See also *Escher on Escher: Exploring the Infinite* by M.C. Escher (1989), *Masters of Deception: Escher, Dali & the Artists of Optical Illusion* by Al Seckel (2007).

Louis Albert Necker (1786-1861), the inventor of the optical illusion now known as the Necker cube, was a Swiss scientist and amateur mountaineer who spent much of his life in Scotland. He is the author of a wide range of books, covering topics as diverse as travelogues in Scotland (*A Voyage to the Hebrides*; *Travels in Scotland*) to birds around Geneva (*Mémoire sure les oiseaux des environs de Genève*, with James David Forbes) to geological studies in the Alps (*Études geologiques dans les Alpes V1*).

Chapter 2a

Medical Applications?

A highly suggestive result with Dutch epilepsy patients

Howard: When did this result come out? When did you first discover this effect?

Diana: I discovered it in the fall of 1973. I still remember the afternoon when it happened, because it really came as a surprise. I then carried out a formal experiment to document the effect, and described the experiment in a talk at the Acoustical Society of America in the Spring of 1974. I also published it as a report in *Nature* in the Fall of that same year[8]. I was also invited by *Scientific American* to write an article about both this and another illusion that stems from it – which I can talk about in a minute – called the Scale Illusion[9].

Howard: So, I'd like to get to the Scale Illusion but, before I do, I have a few fairly basic follow-up questions. Let's suppose that there's a neuroscientist or two sitting next to us right now, what are they going to say about this?

Diana: They're going to say, "It's not true."

Howard: Really?

Diana: Well, a lot of them did, yes.

Howard: Okay, *did*. But now, 40 years on, what would they say?

Diana: Well, nowadays many laboratories have replicated the effect, and some are working on it. In the process they're beginning to understand the pathway that is responsible for pitch perception

8 "An auditory illusion" (*Nature*, 1974, 251(5473), 307-309).

9 "Musical Illusions" (*Scientific American* 233, 92-105 (1975))

and the pathway that is responsible for sound localization in this illusion.

Howard: Right. The question of pathways is interesting to me because when I think about it, it immediately dawns on me that I have no idea what "handedness" actually is, neurophysiologically-speaking. I'm guessing that our brains don't come with little tags on them stating whether or not we are left or right-handed, so I'm thinking that, when you're talking about pathways for pitch and pathways for localization, at one level it seems to me you're talking metaphorically: you're saying that there are these two things going on and, somehow, we have to be doing these two logically independent things simultaneously, with the Octave Illusion being an example where you can mess them up so they don't actually work coherently together.

So that's all metaphorical. But then, presumably, these pathways are somehow real things instantiated in the brain – real physical pathways with real neurons firing.

If I were a neuroscientist, then, I would think that this would be a very captivating result that I'd be dying to understand and explain. After all, it's an empirically justified phenomenon that's been around for 40 years now.

So, what's our best guess at what is actually going on? First of all, how can it be different for people who are different-handed? Is this just one of many effects that are manifested differently for different people depending on their handedness? And where are these pathways in the brain, exactly? Are people working on this?

Diana: Well, I should talk about an experiment that I did in collaboration with a group at the medical school in Utrecht, The Netherlands[10] carried out with my collaboration. There's a test known

10 Brain Center Rudolf Magnus, UMC Utrecht.

as the "Wada test," that you sometimes give to people who have epilepsy that is sufficiently severe that you have to do surgery to control it. But this surgery carries the danger that you might disturb the speech hemisphere (which is the dominant hemisphere) and render the patient aphasic for life.

The Wada test, then, involves injecting an anesthetic into a carotid artery, so as to temporarily anesthetize one hemisphere, and if the patient can still talk, you conclude that it's safe to go ahead with the procedure. Handedness, in itself, correlates very strongly with cerebral dominance, but not absolutely. You get some people, particularly, for example, right-handers with a left-handed sibling or parent, who may very well have an unusual pattern of cerebral dominance. You can't only rely just on the patient's handedness, then, to be sure that you can do the surgery safely.

The Wada test has been done for a very long time now, over half a century, but the problem with it is that it's an invasive test: there are side effects which can sometimes be very serious, and people really don't like to do it unless they absolutely have to. Another possibility is fMRI, where you determine which side of the brain is most active when the person is speaking or listening to speech, but there are questions about whether or not it is as reliable as we'd like. A further possibility is MEG[11], but that is even less reliable than fMRI.

But the statistics in the Octave Illusion concerning the correlation between which ear hears the high tone and handedness (and familial handedness background) is so strong, that we decided that we would study patients at Utrecht who were going to undergo the Wada test anyway, and correlate how they heard the Octave Illusion with how they performed later on the Wada test.

11 MEG is an abbreviation for magnetoencephalography, a neuroimaging technique that involves mapping brain activity by recording magnetic fields produced by the brain.

This was the experiment. First the patients would experience the Octave Illusion and write down what they heard. And then, on another occasion – it could be days or even months later – they would take the Wada test. And in 17 out of 17 subjects, the Octave Illusion correctly predicted the side of cerebral dominance determined by the Wada Test. That's really very exciting, I think. It means that the Octave Illusion might replace the Wada test as a reliable indicator for cerebral dominance, so that patients wouldn't have to go through such an invasive procedure. Or alternatively the Octave Illusion could be used as an adjunct to other tests.

Perhaps the strongest reason that we were encouraged to proceed with this experiment was that, very early on, there was a subject who was completely right-handed, with only right-handers in the family. She seemed, for all the world, to be left-hemisphere dominant. But on the Octave Illusion test, she turned out to hear the high tone on the left – from which one would predict that she was right hemisphere dominant. But when she had the Wada test, she turned out indeed to be right hemisphere dominant – and this was predicted from how she heard the Octave Illusion, but opposite what you would expect of a right-hander.

Further References

The Wada test (also known as the intracarotid sodium amobarbital procedure, or ISAP) is named after neurologist and epileptologist Juhn Atsushi Wada (b. 1924), Professor Emeritus of Psychiatry at the University of British Columbia, who developed the test while a medical resident in Japan in the 1940s.

More formal remarks on the study at Utrecht, which Diana refers to in this chapter, can be found on "Commentary on 'The octave illusion and handedness: A replication of Diana Deutsch's 1974 study'" (*Musicae Scientiae* 17(3) 290-292, 2013). Those interested in more details of the Wada test vis-à-vis fMRI and other non-invasive techniques should see, for example, "Can fMRI safely replace the Wada test for preoperative assessment of language lateralization? A meta-analysis and systemic review" by P.R. Bauer et al. (*J Neurol Neurosurg Psychiatry* 2014 May;85(5):581-8).

Chapter 2b

The Science of Eyes vs. Ears

The neurophysiological differences between vision and hearing

Howard: That's obviously extremely important in terms of the immediate relevance and application of this work, to develop a reliable, non-invasive diagnostic tool for sufferers of extreme epilepsy, and possibly other conditions as well.

But I'd like to return to this question of what we think is actually going on in the brain when we are experiencing these auditory illusions, these two different pathways and processes and so on.

Again, I'm going to return to what you said earlier: that a neuroscientist would just deny this effect.

Diana: Well, initially they did; but you're right, that would never happen now. But at the time, when I talked about it, people just thought that it couldn't possibly be true.

One problem was that it's actually very difficult to set the experiment up properly. What I did (using these Wavetek function generators) was to switch between the low, 400 Hertz tone and the high, 800 Hertz tone and then back again with phase continuity: there were no amplitude jumps at the transitions between the high and low tones, and the phase didn't change direction during a transition. In other words, there was a minimum of clicks in the sound pattern.

But the problem was that people thought this wouldn't matter. So some people tried to set the illusion up by using two free-running oscillators and then using an audio switch to send the high and low tones back and forth between the ears. With this setup you would hear clicks prominently and it just didn't sound that good.

And then there was the fact that there was nothing theoretical, and nothing in any other experimental work that had been done, that would even hint that such an effect would exist. So people were, I think, justifiably, skeptical about it.

Howard: Okay, but now that it's been well established…

Diana: Mostly, I have to say, because I published the CD. I came out with my first CD, *Musical Illusions and Paradoxes*, because I wanted people to be able to hear the sounds as they should be for this illusion.

Howard: And people did pay attention to that?

Diana: Yes. It was there, so you couldn't deny it.

Howard: Getting back to the idea of these two different systems – one for tone recognition and the other for localization…

Diana: I should say that it's true that there are other possible explanations if you don't do further experiments, but I did do another one.

In this experiment, you also have a repeating sequence, but now it's more complicated: Suppose the right ear gets three high tones alternating with two low tones, and the left ear gets three low tones alternating with two high tones. The person who's right-eared *hears* three high tones on the right alternating with two low tones on the left. And when you switch the headphones around, this right-eared person then hears *two* high tones to the right alternating with *three* low tones to the left.

It takes a while to assimilate this because it's so counter-intuitive, but the point is that this experiment showed that people were hearing the pattern of pitches that were presented to one ear rather than the other, and they were localizing each tone to the ear that received the higher frequency. The problem with the Octave Illusion that I think you're implying in your question is that there are other possible explanations because it's a symmetrical pattern.

Howard: So, here you've deliberately made things asymmetrical.

Diana: That's right: three of one sort alternating with two of the other, so you can separate them out. And once again it did show that the "*what*" and the "*where*" systems were separate.

And I showed this in other experiments as well. For example, suppose you have a right-eared person listening to the Octave Illusion – they're hearing the high tone on the right. Then you gradually increase the amplitude of the tone in the left ear, so that, at some point, they're starting to hear the high tone on the left instead, because they're now following the left ear for pitch. And you can measure the point at which that transition takes place. And you can also gradually change the amplitude of the low tone relative to the high tone, so that instead of localizing the tone to the ear that receives the higher frequency, the person will at some point localize to the higher amplitude.

So looking at the results from these manipulations, it turned out that some people showed very strong ear-dominance effects – that is, hearing the high tone on the right or left for pitch – and be very "fussy" as far as localization is concerned: you could change the relative amplitude by a small amount and that would affect their sense of localization. For other subjects, it would be the other way around – they would show strong localization effects but weak ear-dominance effects – measured by the amount of amplitude difference between the ears needed to counteract the illusion. So, there are various ways of showing that the two pathways act in different ways.

Howard: From my recent discussions with other neuroscientists, my sense is that, generally speaking, this view that there are different streams of processing – a "*what*" stream and a "*where*" stream, say – is completely in keeping with our contemporary, neuroscientific understanding.

Diana: Yes, that's true. But, at that time, it was unknown. It's true that some people still argue about it in the case of hearing, but I think that they're now beginning to accept it. It has been known for quite a while now that there are separate pathways (dorsal and ventral) for "*what*" and for "*where*" in the case of vision – but that came later. And at the time that I proposed this for the Octave Illusion, it was not known.

Howard: There seem to be so many neuroscientists who work in vision. As it happens, I recently talked to Kalanit Grill-Spector[12], a vision specialist at Stanford who has identified all sorts of really interesting aspects of facial recognition, and my understanding is that this notion of parallel or different streams of information processing plays a key role in her research. But I'm thinking that auditory processing seems somewhat different: I'm not aware of as many cognitive scientists working in this area. Is that just me, or is there an objective distinction? Are there lots of neuro-scientists who work with that as well? Are there links, perhaps, between the two fields?

Diana: Very few people in hearing have worked on separate "*what*" and "*where*" processes. There is a group in Finland[13] who used the Octave Illusion; they have a hypothesis about where the "*what*" and where the "*where*" components are. According to them, they're both in the auditory cortex, but different aspects of

12 See the Ideas Roadshow issue *Vision and Perception*, with Kalanit Grill-Spector.

13 Both the Cognitive Brain Research Unit, University of Helsinki and the Brain Research Unit of the O.V. Lounasmaa Laboratory, Aalto University have done work on the octave illusion (see references for further details).

the signal seem to be signaling "*what*" and "*where*." But on the whole, there's remarkably little that's been done on this question. I think that's a pity, because there obviously are two pathways.

Howard: You've already mentioned fMRI and MEG and other diagnostic devices. It's easy to imagine, as the technology improves (and it's improving very steadily), all sorts of higher-level experiments to try to isolate these two pathways neurophysiologically.

Diana: Well, a long time ago, Norman Geschwind[14], who was one of the great figures in neuroscience in the middle of the 20th century, wanted to work with me on the Octave Illusion. His idea was to use stroke patients to try to find out where these pathways were.

He said that there were two basic hypotheses. One was that inhibition in a transcallosal pathway between the two hemispheres was involved; and the other was that there were inhibitory interactions in the ascending pathway within one hemisphere. He thought that the Octave Illusion either occurred transcallosally, or it occurred somewhere in the ascending pathway.

He felt that, by studying stroke patients where he knew where the brain damage was, he would be able to shed some light on this question. Unfortunately, we never got around to doing the study because he passed away. But, yes, it should be done; and I suppose I really ought to be collaborating with people and doing this work myself.

Howard: Well, that's not what I meant, of course. I'm not blaming you.

Diana: No, I know. But, of course, you're absolutely right: it does follow, doesn't it? Although I'm not sure whether fMRI, as it stands

---

14 Norman Geschwind (1926-1984), American behavioural neurologist and professor at Harvard Medical School.

presently, would have a sufficiently high resolution to tell us what is happening. I don't know.

Howard: Well, you never know how these things develop, of course, and serendipity often plays a large role. Again, I'm reminded of my conversation with Kalanit Grill-Spector. She was telling me about some seminal results on facial recognition that came from an epilepsy patient. It was a similar sort of thing to what you were talking about earlier – there were some patients with severe epilepsy who came to a local hospital for a series of diagnostic tests to determine whether or not they should have a particular surgical procedure, and sometimes patients agree to participate in additional testing when they are in this diagnostic stage. It was during one of these situations that she happened to discover a spectacular result regarding the localization of neurons responsible for facial recognition.

Diana: Similarly I know that, from my colleagues in Utrecht, when a patient is on the operating table, they do, in fact, stimulate the area that they're thinking of excising, to make doubly sure that it's safe to proceed.

Howard: Right. So, you can certainly imagine doing something like this. Obviously many neuroscientists work in vision, and that's obviously very important, but I'm wondering, listening to you, how many people consider all of this from an auditory perspective – not only in and of itself (which is certainly fascinating enough) but perhaps also involving the conjunction of the two. Maybe if you were to combine vision systems and auditory systems in some way, some other unexpected thing might happen.

Diana: Absolutely, I agree. I think vision has a very long history of looking at high-level processing, whereas hearing, on the other hand, came up through engineering. It *had* to come up through engineering, of course, because that was the only way that you could generate sounds with precision and know what you had.

> So initially most of the auditory researchers were engineers rather than psychologists, and the two histories are quite different. If you look at the history of hearing in the early to middle 20th century – it was really driven by engineering considerations, much less so than vision.

Further References

Some articles on possible neurological models of the Octave Illusion include "Binaural interaction and the octave illusion" by S. Lamminmäki, A. Mandel, L Parkkonen and R. Hari (*J Acoust Soc Am*, 2012 Sep;132(3):1747-53) and "Neural mechanisms of the octave illusion: electrophysiological evidence for central origin" by J. Ross, M. Tervaniemi and R. Näätänen (*Neuroreport* 1996 Dec 20;8(1):303-6).

Norman Geschwind (1926-1984) was a highly influential American behavioral neurologist. For an account of his professional impact, see, for example, "Norman Geschwind (1926-1984)" by Stefano Sandrone (*Journal of Neurology* 260(12): 3197-8, 2013) "Norman Geschwind: Influence on his career and comments on his course on the neurology of behavior" by Orrin Devinsky (*Epilepsy & Behavior* 15(4): 413-6).

Auditory Neuroscience: Making Sense of Sound by Jan Schnupp, Israel Nelken and Andrew King (2011), is a good introduction to our understanding of how we generally perceive sound, as is Andrew Lotto and Lori Holt's 2010 review article, "Psychology of auditory perception" (*Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5) 479-489, Sep/Oct 2011).

Some recent investigations on auditory processes include, "The what, where and how of auditory-object perception" by Jennifer Bizley and Yale Cohen (*Nat Rev Neurosci.* 2013 Oct;14(10):693-707), "Behind the scenes of auditory perception" by S.A. Shamma and C. Micheyl (*Curr Opin Neurobiol.* 2010 Jun;20(3):361-6) and "Brainstem origins for cortical 'what' and 'where' pathways in the auditory system" by N. Kraus and T. Nicol (*Trends Neurosci.* 2005 Apr;28(4):176-181).

Overviews of the science of sound include *The Science of Sound, Third Edition* by Thomas Rossing, F. Richard Moore and Paul Wheeler (2001) and *Physics and Music: The Science of Musical Sound* by Harvey White and Donald White (1980).

Related overviews of the field of psychology of music include *The Psychology of Music (Cognition and Perception), Third Edition* edited by Diana (2013), *Music, Thought, and Feeling: Understanding the Psychology of Music, Second Edition*, by William Forde Thompson (2014), *Psychology of Music: From Sound to Significance* by Siu-Lan Tan, Peter Pfordresher and Rom Harré (2010) and *This Is Your Brain on Music: The Science of a Human Obsession* by Daniel Levitin (2006).

Chapter 3

Sociological Issues

Trusting eyes more than ears, and the benefits of perfect pitch

Howard: Before we move on to the Scale Illusion, I'd like to make a bit of a sociological digression. It's interesting to me that, independent of the details, here's this clearly new and emerging field, and you seem to be stumbling upon one of these things after the other. It's like you're walking among poppies, plucking one after the other. If I were a scientist at the time, I would be saying to myself, "Well, I'm going to go into that field, because this woman is finding all these things."

Diana: I quite agree! It's a mystery to me that there isn't this huge field of auditory illusions, but for some reason there isn't, and I don't quite know why.

Howard: And presumably, there are many things you haven't found yet.

Diana: There must be! This has to just be the tip of the iceberg, so it is rather surprising. I think, partly, people tend to be theoretically-driven, and they also tend to be driven by what they expect, not just by the theory as written in the textbook, but what they really expect to hear. And when they hear something that's so absolutely crazy, such as the Octave Illusion, sometimes they don't like it.

I have to say that I feel that myself too. There's a certain part of me that is made uncomfortable when I have to come to terms with some of these illusions. I don't really want to believe that the world is such a disorganized place. I'd like to feel that when I hear something, I know what I'm hearing and that other people are going to be hearing the same thing, and we can all understand it: one just needs to crunch through the various possibilities and it will all become clear. But when you come across something that's crazy, it's a bit upsetting. And I think that may be partly why not many people investigate auditory illusions.

I've had both experiences. It's exhilarating to know that you've come across something really new, but there's also a sense of discomfort, which is very serious. I think it may often happen that people do come across really fine illusions, and they just think that there's something wrong with them or their equipment, and that they'd better not talk to anybody else about it for fear that all the stuff they'd been publishing in the past has been artifactual or due to faulty equipment.

And it's not only what they think about themselves but, also, what other people will think about them: that they might end up being looked upon as incompetent or something. Those thoughts do enter one's mind. I mean, it's not all that comfortable to stand up and say, "Here's this thing that I claim is happening," and know that some people in the audience are thinking, *Okay, well, it's just a talk, so let's forget about it*.

Howard: Okay, well, I'm a bit confused because I see what that's like for the man or woman on the street: it's very discomfiting to feel that our sense data, to use a philosophical term, might be illusions. Nobody likes to think that, no matter how many degrees she might have.

On the other hand, scientifically, that's one of the most exciting things you can come across: something which is completely counter-intuitive, something that you don't understand and think must result from a flaw in your apparatus or whatever.

However, if you do it over and over again and still get the same result, and then you ask other people to do it and they get the same result, then at some point you realize that you've stumbled upon something which is a completely unexpected, a revolutionary result, for which there is no immediate explanation – and that's about as good as it gets in science. I mean, that's what you're looking for. Hence my analogy about going into this field of poppies. So again, I'm surprised that more scientists aren't completely intrigued by this.

Diana: Well, I think maybe there are two ways to understand this. Partly, I don't mind being wrong, and I think that's important. So, suppose I did an experiment and published a result, and later discovered that I'd been wrong– fine, I can admit that and the world goes on. That's part of it.

The other part of it, I think, is having a pretty good sense of self-confidence about what one is hearing. In listening to the Octave Illusion, I knew that I was hearing a high tone intermittently in one ear and a low tone intermittently in the other ear; that's what I was really hearing and nobody was going to tell me otherwise. But, I think a lot of people, especially for hearing, are not that certain of their own percepts.

Howard: Do you think that most people feel more grounded with vision than they do with hearing?

Diana: Yes. I think people find it difficult to articulate what they're hearing. So, if they hear a mush of stuff that doesn't sound quite right, they can just throw up their hands and blame themselves.

Howard: You spoke earlier about how when you were a small child and couldn't understand why adults couldn't just tell you what note they were hearing – I presume that means that you have perfect pitch.

Diana: Yes. I have had perfect pitch all my life until very recently. It tends to deteriorate as you get older, and people don't really know why. The belief now is that it's something to do with the tissues of the inner ear, but I'm not sure. It's not that bad, but I'm beginning to make mistakes.

Howard: Well, welcome to the world of the rest of us.

Diana: Yes. It's funny how it's upsetting. If you talk to anybody with absolute pitch at the point where they begin to make mistakes, they'll tell you it's very upsetting, and that they don't know why.

Howard: Has your perfect pitch been actually held against you in some way? I can imagine somebody saying, "Here's this auditorily-fixated woman who has these particular skills in this area – she has perfect pitch, she's a music person – and as a result, she's concerned about these things that really don't pertain to the rest of us." Has any of that been thrown your way?

Diana: Not that I've heard. They may be saying that, but nobody's said it directly to me. Although I'm working on perfect pitch a lot now, I'm certainly not fixated on the idea. In fact, it wasn't until fairly recently that I started working on it – because, actually, to have perfect pitch is truly no big deal at all; it's just like seeing a color and naming it. People who have perfect pitch are really puzzled as to why people who don't have it think it's so strange.

Further References

Perfect pitch is discussed in many popular books, including *Musicophilia: Tales of Music and the Brain* by Oliver Sacks (2007), and *This Is Your Brain on Music* by Daniel Levitin (2006).

For a more scientific take on perfect pitch, see, for example, Diana's 2002 article "The puzzle of absolute pitch" (*Current Directions in Psychological Science*, 2002, 11, 200-204) "Absolute pitch: perception, coding and controversies" by Daniel Levitin and Susan Rogers (*Trends Cogn Sci.* 2005 Jan;9(1):26-33) and "Perception of musical intervals by absolute pitch possessors" by Ken'ichi Miyazaki (*Music Perception*, Summer 1992, 9(4) 413-426).

Chapter 3a

Auditory Fudging and Neuroplasticity

The Scale Illusion, auditory scene analysis, and evolutionary factors

Howard: Let's move to the Scale Illusion[15] now.

Diana: That was something that I created based on the Octave Illusion. What happened was that, having discovered the Octave Illusion, I was very excited about it. All night long, I would just imagine notes dancing around. I couldn't get any sleep, they were just sort of flickering around. And in the end I thought, *Okay, supposing I had a series of notes going up a scale while at the same time another series of notes is going down the scale, and the notes from both scales are alternating from ear to ear. By analogy with the Octave Illusion, maybe I will also hear all the high tones in the right ear and all the low tones in the left ear, even though, in point of fact, that's not what's going on.*

Fig. 5 (Next page) - The pattern that produces Deutsch's Scale Illusion (a) and a way it is often perceived when listening through headphones (c). The notation in (b) shows how the pattern is composed of ascending and descending scales.

15 See chapter references for links to explanatory notes and MP3 audio file of the scale illusion.

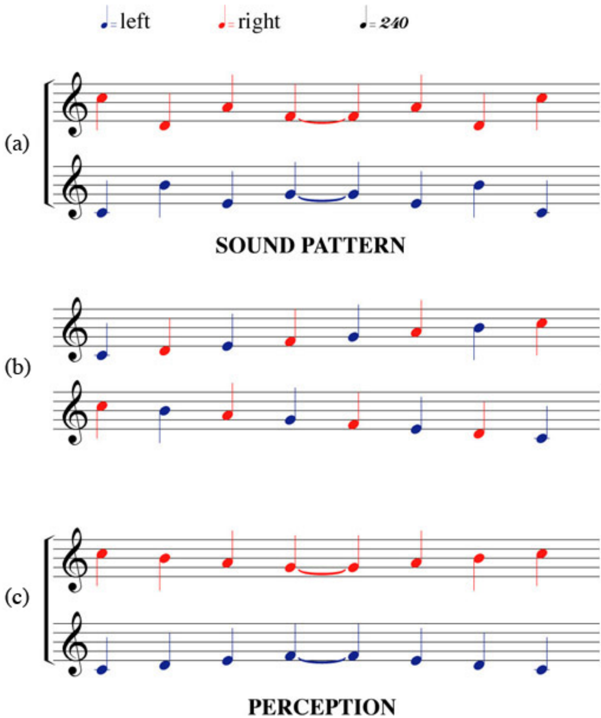

So, early the next morning, I rushed to the lab and typed in the Scale Illusion; and there it was! It was only later that I realized it had to have an entirely different basis, because in this case, you're hearing both sequences of tones at the same time, you're just mislocalizing half of them.

Howard: So, when you say "a different basis," it's still playing with this idea of these two processing streams, right?

Diana: I don't think so. I think it really has to be different because, in the case of the Octave Illusion, you're fusing the two tones that are separated by an octave into a single tone. However, in the case of the scale illusion, you continue to hear all the tones; it's just that half of them are perceptually displaced to the opposite ear. I think the scale illusion occurs in part because the idea that you'd have two sound sources in different positions in space that are each producing tones that leap around in pitch is just improbable. On the idea of "unconscious inference", as Helmholtz[16] used to describe it – it's just improbable. So instead, we opt for the more probable interpretation.

You should really listen to it through headphones because, again, most right-handers hear the higher tones on the right and the lower tones on the left – but this won't work if you don't use headphones. However, listening through loudspeakers you can still hear two clearly distinct melodies – a high one going smoothly up and then down with a simultaneous low one going smoothly down then up, with the two of them meeting in the middle.

Now, if I mute one of the channels, what you find is a melody that jumps around and doesn't move smoothly. On just one channel, then, you hear an entirely different melody than what you hear when both channels are sounding. There's a better, more dramatic example that I have called "The Chromatic Illusion[17]," where you have a two-octave, chromatic scale. It's the same principle but, because they are so separated from each other at the extreme and because they're moving in semitone steps the whole way, it sounds more dramatic.

16 Hermann von Helmoltz (1821-1894), German physicist and polymath. See chapter references for more details.

17 See chapter references for links to explanatory notes and MP3 audio file of the chromatic illusion.

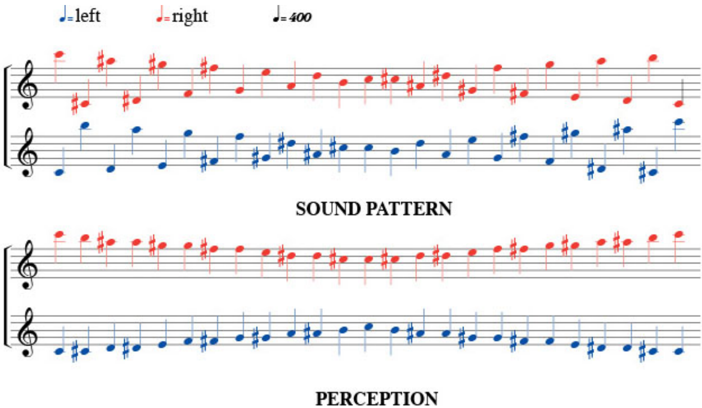

Fig. 6 - The pattern that produces Deutsch's Chromatic Illusion, and a way it is often perceived.

© 2015 Diana Deutsch. Reproduced with permission

Again, in this example when both channels are producing sound, you hear two distinct melodies moving smoothly in pitch. But when you isolate either of the channels, you hear a completely different melody that jumps around in pitch. This illusion is more accessible, I think, because what you hear when you're listening to each channel separately is dramatically different from when you hear the two together. That's even true when you're not listening through headphones, although it's a much more dramatic effect with the headphones.

Howard: Tell me more about what you mean by "unconscious inference" and what you think is going on here.

Diana: Well, I think that the auditory system is a messy one. If we think of it as beginning at the front end, there are something like about 125 million receptors in each eye and only 16,000 in each ear; and only about 3,500 to 4,000 of those in each ear are afferents, meaning that they actually send information up to the brain.

So when you stack the numbers up, you can clearly see how dramatically low the number of receptors assigned to your ears is, especially given that the information they have to deal with is incredibly impoverished. Then if you consider the amount of brain tissue that's involved in vision as opposed to hearing, the difference is beyond enormous. In fact, if you think about it, it's really rather amazing that we're able to use our auditory systems at all to make much sense of the world.

The auditory system is used primarily for communication – using speech and music – and it's very highly specialized for that purpose. (It's very good with rapid timing and so forth, which is needed for speech and music). But I don't think that the auditory system is much good at scene analysis, which is what we're talking about here – a sound source on one side of the listener and another sound source on the other side. The auditory system just isn't able to grasp what's really going on if you take the whole thing from the bottom-up. It uses what we know is probable based on past experience – maybe some aspects of what we know are also inbuilt – but the result is that we have to, somehow, fudge our way through what we're hearing when it comes to hearing auditory scenes.

Howard: This disparity of dedicated resources for the two information-processing systems – auditory versus visual – is shocking to me; I simply had no idea.

One hears about the plasticity of the brain and about how different systems that lie dormant for extended periods of time might change. So I'm wondering if you've thought of conducting experiments with blind people who might have had some "reordering", as it were, of their neurons, with some of these areas that would have otherwise been devoted to vision now rededicated to "ramp up" their auditory system's development and processing. Perhaps their responses might be somewhat different than those of other people.

Diana: Well, that does seem to be the case. So, for example, the prevalence of absolute or "perfect" pitch in blind people is very high.

Howard: Really? So, it's statistically significant?

Diana: It's huge, yes; very high indeed.

Howard: Does it matter whether they're blind from birth or became blind later on?

Diana: That's an interesting question. I can't tell you for sure. But for people who are blind from birth or from a very early age, it's definitely true. And in addition to absolute pitch, it's also pretty well known that people who are blind are able to tell a lot about spatial aspects of sounds that they hear.

There hasn't been an awful lot of work on auditory perception in blind people to the best of my knowledge, apart from those two aspects I was talking about, but I bet it's true that people who are blind have a far better auditory system.

But on the whole, for understanding what's going on in the world and being able to locate objects in the world, vision dwarfs anything that can be done with hearing. People have tried to do auditory scene analysis but the results are pitiful compared to what you can do with vision.

Howard: Again, I'm struck by the potential of using aspects of your work as a window, as another tool, towards more deeply examining some of these neurophysiological processes, particularly in light of recent advances in neuroscience.

Diana: There's another problem when it comes to the neuroscience of hearing in humans. My understanding (though I'm not an expert at this) is that rhesus monkeys are pretty similar to humans in vision, but it's not like that when it comes to hearing, because monkeys don't talk or come close to talking. They have monkey calls and so forth, but they don't have the hierarchal structure of sentences that people have.

Howard: Sure, but that makes it very interesting, of course, because one naturally starts to consider things from an evolutionary perspective: what distinguishes us, what separates us from these animals.

Diana: Indeed, but it does make it very difficult to, say, carry out physiological experiments on animals and generalize the results to humans. Now, there have been some very interesting experiments which have been done, particularly on what's known as "stream segregation", where rapid sequences of tones alternate between two different frequency ranges, and if the frequency ranges are very close in pitch, people will hear them as a single, coherent sequence, but as the frequency ranges diverge, they'll hear two disconnected sequences in parallel. People have shown very convincingly, at the neurophysiological level, how this comes about.

So there are some things that can be done, but there are limitations too. So in defense of people who are doing auditory neurophysiology on animals and not trying to relate their findings to human perception, there are limits to the types of questions that they can ask and assume that their findings will hold for people – at least at our present stage of technology.

Further References

See Diana's dedicated webpage on the Scale Illusion for added explanations, MP3 audio files, explanatory videos and references.

See Diana's dedicated webpage on the Chromatic Illusion for added explanations, MP3 audio file and reference.

Hermann von Helmholtz (1821-1894) was a remarkably diverse thinker who made seminal contributions to physics, philosophy, psychology and physiology. His theory of unconscious inference (for vision) was explicitly set out in his *Treatise on Physiological Optics*. A good introduction to this topic is the article "Perception as Unconscious Inference" by Gary Hatfield, Chapter 5 of the book *Perception and the Physical World: Psychological and Philosophical Issues in Perception*, edited by Dieter Heyer and Rainer Mausfeld (2002).

A book by Helmholtz which is relevant to this conversation is *On the sensations of tone as a physiological basis for the theory of music*. Books about Helmholtz and music (among other things) include *Helmholtz and the Modern Listener* by Benjamin Steege (2012) and *Helmholtz: From Enlightenment to Neuroscience* by Michel Meulders (2010).

For those interested in more details of auditory scene analysis, see, for example, the 2014 review articles "Probing auditory scene analysis" by S. Deike, S. L. Denham and E. Sussman (*Front Neurosci.* 2014 Sep 12;8;293) or "Predictability effects in auditory scene analysis: a review" by A. Bendixen (*Front Neurosci.* 2014 Mar 31;8;60). A classic reference is, *Auditory Scene Analysis: The Perceptual Organization of Sound* by Alfred Bregman (1990).

Scientific work on human echolocation includes "The size-weight illusion induced through human echolocation" by Gavin Buckingham, Jennifer Milne, Caitlin Byrne and Melvyn Goodale (*Psychological Science*, Feb 2015, 26(2):237-242). See also *Beginner's Guide to Echolocation for the Blind and Visually Impaired: Learning to See With Your Ears* by Tim Johnson (2012).

More information on auditory stream segregation can be found in, for example, *Auditory Cognition and Human Performance: Research and Applications* by Carryl Baldwin (2012) and *Sensation and Perception, Ninth Edition* by E. Bruce Goldstein (2014).

Technical papers on auditory stream segregation include "Auditory stream segregation and the perception of across-frequency synchrony" by Christophe Micheyl, Cynthia Hunter and Andrew Oxenham (*J Exp Psychol Hum Percept Perform.* 2010 Aug;36(4): 1029-1039), "Towards a neurophysiological theory of auditory stream segregation" by J.S. Snyder and C. Alain (*Psych Bull.* 2007 Sep;133(5):780-99) and "ARTSTREAM: a neural network model of auditory scene analysis and source segregation" by S. Grossberg, K.K. Govindarajan, L.L. Wyse and M.A. Cohen (*Neural Netw.* 2004 May;17(4):511-36).

Chapter 3b

Surrounded by Illusions

From the Glissando Illusion to Tchaikovsky's 6th

Diana: There's another illusion, where you have an oboe tone that alternates with a glissando that goes up and down in pitch and switches back and forth. You really do need to listen to this through loudspeakers rather than through headphones. What people hear is the oboe tone going back and forth very clearly, but the sine wave appears to be moving around in space in accordance with its pitch motion. It's a very striking effect.

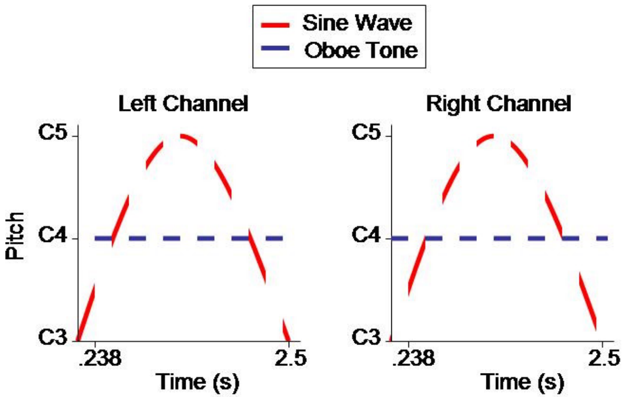

Fig. 7 - Fragment of the pattern that gives rise to Deutsch's Glissando Illusion, as it was presented in the experiment.

© 2015 Diana Deutsch. Reproduced with permission

In some ways it's more striking because you don't hear bits of the glissando[18] at all, you just hear a completely seamless glissando, even though what's being played are bits and pieces of one.

I should also mention that this is an example that composers really like, because if you're in an auditorium that has sufficient reverb, it really sounds like there's this frisbee or something that seems to start off to your left and move gradually to your right as its pitch increases – at its highest pitch, it appears to sit up in the right hand corner – and then it comes back down again. This has a particular aesthetic.

Howard: Different people have different reactions?

Diana: There is, in fact, a handedness correlation with the glissando illusion. If you're strongly right-handed, you tend to hear the highest pitch as the highest in space and off to the right, but if you're not strongly right-handed, you can instead get a lot of different percepts. You don't even have to be left-handed, you just have to not be strongly right-handed for those statistics to kick in.

Basically you sit there listening, and the glissando appears to be moving around. If this is at a concert, you're told in the program notes that "This isn't really happening," and you enjoy your own very strange percept. So that's a good one to play in an auditorium.

Howard: And what would the scientific interpretation of this be?

Diana: Again, I would say again that the interpretation here would in terms of unconscious inference. Just as in the Scale Illusion and the Chromatic Illusion, it's very improbable that there's a source that has a section of a sine wave coming from one position in

18 A glissando is a musical term, referring to a glide from one pitch to another.

space and then suddenly, it turns into a section of an oboe tone and then it's a section of a sine wave again. And so on. It's much more probable, in terms of the real world, that you have one source that's producing a sine wave that's moving slowly through space and another source that's producing an oboe tone.

Actually, in this case, you do hear the oboe tone correctly as going back and forth; but I think that's because it's a very rich tone – it's rich in harmonics – so the localization cues are very strong. But for the sine wave the localization cues are obviously very weak, because it's only just a sine wave that you hear.

Howard: I don't understand that. So, when you talk about the different cues and a rich tone with lots of harmonics, the idea is that it's easier to localize something that has more harmonics?

Diana: Yes, that's something that's well known from other literature. However, people are not very good at localizing just a sine wave, so in this case, they're just making a bet as to where it is coming from. Again, I don't think it has to do with separate "*what*" and "*where*" channels as much as the fact that the auditory system is trying to do scene-analysis: it's presented with something that's very highly improbable and it discards the highly improbable but correct interpretation and, instead, opts for something which is much more likely, even though it's incorrect.

Now, in the case of vision, there are some things like that as well. For example, there's this famous hollow mask picture taken from the inside where you see it as facing outwards, even though, in point of fact, the nose and the lips and so on are all facing inwards. However, you see it as facing outwards because our knowledge and experience of faces is such that they face outwards, so we opt for an incorrect but much more plausible interpretation. There are some other cases in vision where this happens, but it seems to be happening a great deal in hearing.

Fig. 8 - Picture of a hollow mold of a face. We perceive the features of the face to be projecting outward, even though they are projecting inward. This, as the Scale Illusion, provides an example of the strong influence of unconscious inference in perception.

Howard: When you say that it's happening all the time, that brings up an obvious question. It's happening in these particular circumstances that you have arranged but, as you said earlier, it might be happening much more often than we realize. You're conducting very carefully prepared experiments that highlight these phenomena, but they might be happening all the time, all over the place.

Diana: Yes. And in point of fact, the Scale Illusion is a very good example. There's an article that came out in *Scientific American* written by Shawn Carlson[19], in the Amateur Scientist column. I had told him about this story which is pretty well authenticated about an argument between Tchaikovsky and conductor Arthur Nikisch. Tchaikovsky was about to premier his 6th Symphony, in which the final movement begins with a very peculiar passage in which the theme and the accompaniment waft back and forth between the violin parts. The audience hears the theme as coming from one set of violins and the accompaniment as coming from the other; and this is true, even with the orchestra arranged in 19th-century fashion, with the first violins to the left and the

19 "Dissecting the Brain with Sound", *Scientific American*, December 1996.

second violins to the right.

Now, there's no evidence that either Tchaikovsky or Nikisch realized that there was an illusion going on, but they did have an argument about that passage when they came to discuss the premier of the symphony. Nikisch tried to insist that it be re-scored so that the notes from the theme would all be coming from one set of violins and those from the accompaniment from the other. But Tchaikovsky adamantly refused to change the scoring, and the piece was premiered as he had originally written it.

A few conductors still adhere to Nikisch's school of performing of this symphony but, in point of fact, if you play it with the orchestra arranged in 19th-century fashion so that one set of violins plays a jagged sequence and the other set of violins plays a different jagged sequence, you really do hear the whole pattern reorganize itself so that you hear this melody that neither violin part is actually playing.

Nobody knows why Tchaikovsky wanted to score the piece that way and they don't know why Nikisch objected. Some people think that maybe Nikisch, being a conductor and standing between the two violins, really didn't get an illusion, whereas, Tchaikovsky, just imagining it, thought that the audience would hear the theme and accompaniment as separate melodic lines. Another suggestion is that, if you re-score it as Nikisch had wanted, the two parts are easier to play because the notes follow each other in small steps. It's not known; but this is an example of an illusion that is like the Scale Illusion, and that happens out there in the real world. The conductor of the UCSD symphony, Tom Nee, did, in fact, set this up with the orchestra arranged in 19th-century fashion when NOVA came to film this segment for their program, *What is Music?*[20]

---

20 Nova: Season 17, Episode 8, *What Is Music?* (21 November, 1989)

Howard: So, the idea is that the conductor, being in between the two violin sections, would presumably hear things quite differently.

Diana: One would think so, yes. But one can't be sure. There are other examples of how one can generate these sorts of things in non-laboratory settings. I have a video[21] of the Scale Illusion performed and experienced by fifth graders at Atwater School in Shorewood, Wisconsin (created by their teacher, Walt Boyer). It's a fantastic example of how the Scale Illusion can work in an incredibly messy situation.

It's funny; I began by generating examples in the lab, thinking, "Okay, I've got to have sine-wave tones, I've got to have earphones, I've got to be very, very accurate about the transitions," and so on, but it turns out that you can get similar effects when you just have a mesh of violins on one end of a concert hall and another mesh of violins on the other. That was a surprise to me.

Howard: Right. And presumably, not only can you replicate or duplicate these effects under far more mundane or ordinary circumstances, but, as you were saying before, there may be effects that are happening that are completely coincidental, that are just out there that we're not even deliberately trying to engineer at all.

Diana: That's right. But if it weren't for being able to do it with very strict control of the stimulus parameters in the first place, I wouldn't have been able to make that statement about the Scale Illusion. Had I just, say, pointed at the Tchaikovsky piece or simply had a bunch of violinists playing the Scale Illusion, some listeners would have said "Well, it's probably something to do with the harmonics, who knows?"

But I was able to be very precise about it using sine waves and headphones, so there's no other interpretation other than people's brains just switching tones around in space.

---

21 See here for the page containing the video.

Further References

For additional background on content covered in this chapter, see Diana's 2010 *Physics Today* article, "Hearing music in ensembles" (*Physics Today*, 2010, February, 40-45).

For a host of more general examples of auditory illusions and their pervasiveness in everyday life, see, for example, University of Wisconsin music theorist and electrical engineering professor Bill Sethares' webpage, Illusions of Sound Perception, Meara O'Reilly's website and Vimeo page and Tumblr page, Robinson Meyer's 2014 article in *The Atlantic*, "The Illusion That Makes It Sound Like a Pitch Is Constantly Rising", Ellie Zolfagharifard's 2014 online article in The Daily Mail, "Forget optical illusions – these incredible tricks will make your head spin", *Science Daily's* November 2009 article, "Auditory illusion: How our brains can fill in the gap to produce continuous sound" or the YouTube video from AsapSCIENCE, "Can you trust your ears?".

Chapter 4

Perfect Pitch & Tone Languages

Why Mandarin might help your musicianship

Howard: You mentioned perfect pitch earlier. I'd like to talk a little bit more about what you've been able to discover there.

Diana: Well, I first got interested in perfect pitch when I was studying a group of Vietnamese subjects on an entirely different illusion called "The Tritone Paradox[22]". I became very interested in their speech, which sounded very melodious and different from the way we speak.

I found that, when I tried to repeat back a word that one of them had spoken, they had no idea what I was trying to say. They thought either that I was talking nonsense or that I was trying to say something entirely different. I figured that the confusion must be pitch-related.

I'd start off pronouncing a Vietnamese word with a very low pitch and this would make no sense to them. Then I would go a little higher, and it still would make no sense, until finally I'd go a bit higher and then they would understand what I was saying. The word would continue making sense for maybe a semitone or two, and then I'd hit another zone where it would be meaningless. Then I'd suddenly hit another zone where the word would take on an entirely different meaning. I think that happened three times. It appears very odd to an English speaker for a word to have an entirely different meaning depending on its pitch. It seemed to me that this was an example of perfect pitch, because my Vietnamese subjects were so fussy about the exact pitch with which a word was pronounced.

22 See Diana's webpage dedicated to the Tritone Paradox for more information.

At the time, I wasn't able to transpose words using computer software (nowadays, with software, you can transpose anything), but I was able to analyze pitches. So I had the subjects recite a list of ten words on one day, and then the same list of ten words on another day. Then, together with my colleagues Trevor Henthorn[23] and Mark Dolson[24], I analyzed the average pitch of each word. We found[25] that the pitches of the words were remarkably similar on the two days. I think a third of the subjects produced words that, comparing across days, were on the average less than a quarter of a semitone away from each other.

Then I tried the same thing with Mandarin tones, and with a group of Mandarin speakers as subjects. Once again, one third of the subjects were within a quarter of a semitone in terms of pitch consistency when they recited the same list of words on different days. That was remarkable to me.

I gave a talk about this at the Acoustical Society with the rather bold title, "Tone-Language Speakers Have Perfect Pitch", because, it seemed to me that that's what was going on. Of course, people said to me, "Well, okay, it's true for speech; but how do you know they have perfect pitch for music?"

Howard: Well, you can do experiments, presumably.

Diana: Yes, so I did[26]. I thought, "Well, okay, I'll take a group of

23 Trevor Henthorn, Department of Music, UCSD.

24 Mark Dolson, Creative Advanced Technology Laboratory.

25 "Speech patterns heard early in life influence later perception of the tritone paradox" (Deutsch, D., Henthorn, T., Dolson, M., *Music Perception*, Spring 2004, 21(3), 357-372).

26 "Absolute pitch among American and Chinese conservatory students: prevalence differences, and evidence for a speech-related critical period" by Diana, Trevor Henthorn, Elizabeth Marvin and HongShuai Xu (*J Acoust Soc Am* 119(2), 719-722, 2006).

Mandarin speakers and a group of English speakers and look at the prevalence of perfect pitch in these two groups." I ran into a problem before I began, because I had difficulty finding collaborators. The people who spoke English said that the experiment would be a waste of time because there were too few people with perfect pitch, so we would never get significant results, while the people in China said that the experiment would be a waste of time because everybody knows that any musician worth his salt has perfect pitch.

This went on for a couple of years, but I finally found two excellent people to run the experiment: the Dean of the School of Music at Eastman[27] and a graduate student at the Capital Normal University in Beijing[28]. We studied two large groups of students at distinguished music conservatories. The first group were at the Central Conservatory of Music in Beijing – they all spoke Mandarin. The second group were at Eastman School of Music in Rochester, New York – these all spoke English or a different non-tone language.

Fig. 9 - Percentages of subjects who obtained a score of at least 85% correct on the test for absolute pitch, allowing for semitone errors. CCOM: students at the Central Conservatory of Music, Beijing, China; all speakers of Mandarin. ESM: students at Eastman School of Music, Rochester, U.S.A; all non-tone language speakers.

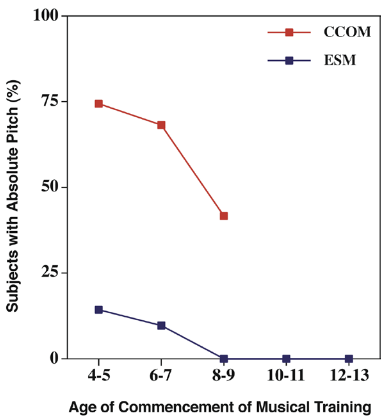

27 Elizabeth Marvin.

28 HongShuai Xu.

The data showed that there was a very strong effect of age of onset of musical training: those who began musical training very early did much better than those who began training late. In fact, as far as the English speakers were concerned, there weren't any with perfect pitch who began musical training at age 8 or 9. (We didn't analyze the data from the few people in the study who began musical training even later than this.)

And when one allows for semitone errors, the Beijing people do even better.

Fig. 10 - Average percentage correct on the test for absolute pitch among students at the Shanghai Conservatory of Music.

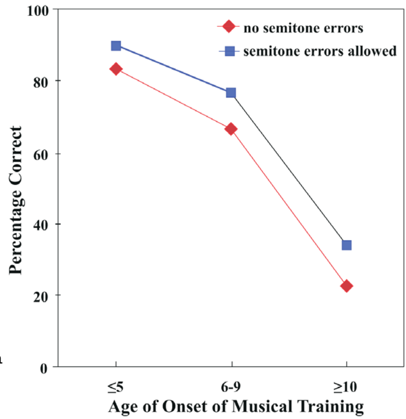

Howard: One thing that would be interesting for me to know – maybe you've already done studies on this or have been thinking about doing studies about this – is to see how the data are correlated with respect to later language learners.

Again, I'm thinking about ideas of neuroplasticity and so forth: if I'm somebody who's learning Chinese later in life, as opposed to a native speaker, if I'm somebody who is suddenly forced to use my neurons in a particular way to focus on these tones and correlate that with my sense of pitch – I imagine that might statistically have some measurable effect on other aspects of my musical ability in terms of being able to isolate pitch.

The thinking is that, if I were born and raised in Oklahoma, say, and then when I'm 10 years old I suddenly move to Beijing and now have to learn Mandarin – which means that an integral part of my existence is now focused on paying attention to these tones, one might (statistically, at least) find that this is also having a secondary effect in terms of my musical development, in terms of sensitivity to pitch.

Diana: I've never heard that suggested, but it makes an awful lot of sense to me. I bet you're right, but I would also imagine that the effect would depend on the age of the child when he or she begins to learn Mandarin.

Howard: It would be a more pronounced effect, presumably, the younger the child would be.

Diana: Yes. I would say that, for children in their teens, they might not do any better than adults. What I think is going on is that absolute pitch is really a feature of speech and that the critical period for speech should absorb the critical period for acquiring absolute pitch, because the time frames involved are so similar.

After all, there is this difference between Mandarin and Cantonese speakers on the one hand and speakers of English and other non-tonal languages on the other. The effect is very pronounced indeed.

Further References

For more information and more recent work on language and pitch, see Diana's article "Absolute pitch" in *The Psychology of Music*, 3rd Edition, (2013). "Absolute pitch among students in an American music conservatory: Association with tone language fluency" *Journal of the Acoustical Society of America*, 2009, 125, 2398-2403.

Additional background can be found in Hazel Muir's 2009 *NewScientist* article, "Tonal languages are the key to perfect pitch", Don Monroe's 2004 *Scientific American* article, "Speaking Tonal Languages Promotes Perfect Pitch" or James Glanz's 1999 *New York Times* article, "Study Links Perfect Pitch to Tonal Language". See (hear) also the *Radiolab* piece, "Sometimes behave so strangely".

Chapter 5

Towards Monotony?

The tonal implications of globalization

Howard: Do you think that there is an evolutionary argument that supports this? If the time for development of absolute pitch is so strongly correlated with the time for development of language, then one might think that there are deeper, evolutionary arguments at play with respect to music and speech.

Diana: I would say that, yes. In fact, that's what I'm expecting to put in my next book. I do think that music and speech probably evolved from a proto-language that included absolute pitch, and that tone language is closer to that earlier language. Quite frankly, I think that non-tone languages such as English are in a sense degraded, in that they don't have tone involved, and I think that's a pity.

I also think that it's a problem for language that there's so much interaction between people who speak different languages – people who speak non-tone languages coming into China, even people who speak different dialects in China interacting with each other. For this reason, I think that this absolute pitch component is going to decline. It's going to have to decline to get to the lowest common denominator. For example, Shanghainese, I believe, used to be a very pronounced tone language, but now people in Shanghai speak what almost sounds like a non-tone language.

Howard: Really? So, what are the forces at play?

Diana: Well, if we assume that, within a linguistic community, tone language involves perfect pitch, then there has to be some commonality of agreement as to what information is conveyed by tones in particular pitches.

We've done specific studies on this as well[29]. We examined two, remote, Chinese villages that are only forty miles from each other, but they're in a very mountainous region, so it takes a long time to get from one of them to the other. Now, the dialects that are spoken by people in the two villages are in the same family of Mandarin dialect, but they are still somewhat different.

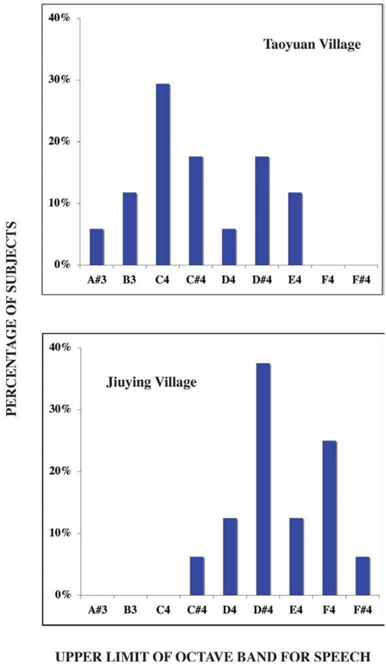

Fig. 11 - Pitch ranges of speech in the two villages at the border of Hubei and Chongqing. Taoyuan Village is in Hubei, and Juiying Village is in Chongqing. The graphs show the percentages of subjects for whom the upper limit of the octave band for speech fell in each semitone bin.

We had people in the two villages recite a three-minute passage, and every five milliseconds we took pitch samples of their speech, and we found that the pitch ranges in the two villages were very different.

29 "The pitch levels of female speech in two Chinese villages" (Deutsch, D., Le, J., Shen, J., Henthorn, T., *J Acoust Soc Am* 2009 May;125(5):EL208-13).

Our thinking was that, in isolated villages, people are going to speak in roughly the same pitch range as each other. The fact of having the same pitch range helps them to gage very rapidly the emotional tone of the speaker, as well as to apprehend the grammatical structure of an utterance quickly, and to know what verbal meaning is intended (since tone language is being spoken).

Now, supposing you have someone who comes to the village speaking a different dialect. If the pitch range in which he speaks is different from the range that is generally used in the village, people could experience difficulty understanding him. So dropping the importance of pitch range in speech makes it easier for people from different regions to understand each other.

Howard: I see. So, that's the force that leads to a more monotone nature of these languages?

Diana: Yes, that's what I think is happening. In Japan, for example, you have pitch accent, which is like a degraded form of tone language. With pitch accent, you have two syllables and, depending upon the pitch relationship between the syllables, you sometimes have different meanings.

There's an example where the words for 'bridge' and 'chopsticks' involve pitch differences between two syllables. For example, in Tokyo, the word 'hashi' spoken as 'high-low' means 'chopsticks; spoken in low-high it means 'bridge', and spoken on the same pitch it means 'edge'. But in Kyoto the opposite happens: 'hashi' spoken high-low means 'bridge', and spoken low-high means 'chopsticks'. So you can see that, ultimately, they're just going to drop this aspect of speech, since it can lead to confusion when people who speak different dialects are communicating.

Howard: This makes me think of something else that you said earlier. I'm thinking of the issue of "musical inference" that was highlighted with respect to some of these illusions you showed me earlier, like the Scale Illusion and the Glissando Illusion.

Perhaps if you're somebody who lives in a place like Japan or China where you're relying upon these informational processes more than non-tonal language speakers, you might imagine that these illusions which are predicated on faulty inferences would be highlighted in some way – maybe stronger, or weaker, I don't know. Do you see where I'm going here?

Diana: Yes. And, in fact, that's right. Let's take Mandarin. The same tone has multiple meanings for homophones, and you have to do an enormous amount of inference to understand exactly what people are saying.

Howard: Right. So, have you done comparative studies with the Glissando Illusion where you compare doing it in Tokyo and Shanghai with Oklahoma, say, and see if there's a greater or lesser effect because of the different average levels of inference sensitivity due to tone language training (or not)?

Diana: I haven't, but I bet you're right. I never thought of that specifically, but I bet it's correct. There was a linguist in the early part of the 20th century[30] who, in order to make the argument that Chinese shouldn't be using what they call Pinyin (the phonetic system for transcribing Mandarin into the Latin alphabet) and should instead stick to the original Chinese characters, composed a story which consisted of about a hundred words about a poet who eats lions[31]; and it's just the word "shi" for the whole time. It's a perfectly understandable, coherent and reasonable story, all on the word "shi". Obviously, this means that there have to be

30 Yuen Ren Chao (1892-1982), Chinese-American linguist and educator.

31 Lion-Eating Poet in the Stone Den – see here for a 37-second YouTube video.

several different meanings even within one tone.

Just to emphasize your point, then, I think you're absolutely right; there's an enormous amount of inference that has to go with all of these tone languages. People really need context in order to understand what's actually being said.

Howard: And presumably the music that people from this part of the world have created is very different and somehow related to this tonal structure. Maybe "structure" is the wrong word here, but one would imagine that their particular musical heritage would somehow be linked to the tonal nature of their language.

Diana: Yes. In fact, I'm doing an experiment right now with a graduate student in Guangzhou where we're looking at memory for pitch using Cantonese tones. (In Guangzhou people speak Cantonese as well as Mandarin, and many are more familiar with Cantonese.)

I think that it's really interesting to consider the effect of the meaning of these tones on short term memory for pitches. This graduate student told me that, when attempting this short term memory task, and the tones are in the pitch range for speech, she finds it very difficult, because she keeps on thinking of the meaning of the words. (But she does well regardless).

When Chinese comedians perform in music halls, one of the things they'll sometimes do is play a phrase on an erhu, say – which is a stringed instrument somewhat like a cello – and the audience will understand the intended verbal meaning, and they find this incredibly funny.

I know that one example that works really well is the word for "hello", but I really want to get a full phrase that works. This shows how tremendously important the musical nature of speech is for tone-language speakers.

Howard: Talking to you, I almost get the sense that there is some regret that your native language wasn't a tone language.

Diana: Yes, you're right. I have tremendous respect for tone language. I think it's fantastic, and it sounds so beautiful too.

Further References

Yuen Ren Chao (1892-1982) was an extremely accomplished Chinese-American scholar. Born in China, he earned a PhD in philosophy from Harvard University, and later served as Bertrand Russell's interpreter when the English philosopher visited China in 1920. He later became the Agassiz Professor of Oriental Languages at UC Berkeley. His many published works include *A Grammar of Spoken Chinese* (1965), *Language and Symbolic Systems* (1968), an autobiography, *Life with Chaos* (1975) and a highly recognized translation of Lewis Carroll's *Alice's Adventures in Wonderland* (*Alisi manyou qijing ji*, 1922).

Chapter 6

Embracing Discomfort

The Speech-to-Song Illusion and the benefits of being confused

Howard: Let me ask you a different sort of question, now. Pretend for a moment that I am an omniscient being and I could answer any sort of scientific question that you'd like to pose: the nature of music, the nature of language, the relationship between the two, what's happening neurophysiologically, what's happening elsewhere, and so forth. What sorts of questions would you ask me? What are you dying to know?

Diana: Well, my main thrust would involve the question of illusion and reality, which gets me back to the work that I did way back in Oxford as an undergraduate when I was studying philosophy. I've always been interested in illusion and reality. Perhaps it means that my sense of reality isn't as robust as that of other people, but that is a natural focus for me.

Howard: Have you made progress with respect to distinguishing between illusion and reality, or are you just as (or more?) confused as you were back at Oxford?

Diana: Well, I'm even more convinced than I was that what we consider to be reality may be illusion. I'm speaking very broadly in order to try to answer your question.

Howard: That's great. I'd like you to speak as broadly as possible.

Diana: So, yes. I feel that we really can't necessarily trust our senses, and this may very well expand into other things as well.

Howard: By "other things," do you mean other senses? Other aspects of perception?

Diana: Other senses, primarily, yes. But I suppose I'm very agnostic indeed about reality. I've become more and more agnostic as the years have gone by, mostly because I've been constantly surprised in my research. If I approach something like a sound pattern, I'm no longer making bets about what I'm going to hear. I'm not sure if that really quite answers your question. But yes, I've always been concerned about this whole issue of how much is illusion and how much is reality. I don't think I just stumbled on these results completely by accident; these issues are meaningful to me.

Howard: Has all of this work changed your appreciation of music personally? Do you listen to or play music in a different way? Do you get enjoyment from it in a way that you might not have otherwise had you not been experiencing all of these effects?

Diana: I would say probably not, actually. I think I respond to music aesthetically, and it's rather disassociated from all of this.

Howard: It's more of an aesthetic, personal choice as opposed to, as you said, this "window" into illusion versus reality? That might be one good way to phrase it: that all of your musical research is exploring this boundary between illusion and reality?

Diana: Yes. One thing we didn't get a chance to talk about is the "Speech-to-Song" illusion, which I know people are fond of hearing.

I was fine-tuning the spoken commentary on my first CD, *Musical Illusions and Paradoxes*; and when you do post-production in audio, you have to do things like "softening" S's and "unpopping" P's, and so on.

I was working on this particular phrase, which happened to be, "Sometimes behave so strangely," and at the time, I had been doing a lot of work, and I was kind of tired and so I began working on something else and forgot about it.

Suddenly, it seemed to me that a strange woman had come into the room and had started singing. I didn't at first recognize that my own voice was being spewed out of the loudspeakers, but instead it sounded absolutely like song and it didn't sound at all like the way it had sounded before. But at all events, I again collected myself and decided to work on the effect.

This is how it works. First, you hear the entire sentence by itself, and then you hear the specific part of the phrase being looped. First it continues to sound like speech, but after a while you start to hear it being sung rather than spoken.

Howard: It sounds like there's a discontinuity when you go back and listen to the sentence again – it's very different from before.

Diana: Yes. In point of fact, one of the interesting things about this is that, once you've heard it, your brain's been altered for that phrase, spoken in that way, so that a year or two can go by but you'll still continue to hear it as song. I don't understand it fully, but it's a clear example of a very rapid and extremely long-lasting case of neuroplasticity.

Howard: If you were to talk to a neuroscientist about this and show them this "Speech-to-Song" example, what would the reaction be? Would it be, "That's interesting," or, "I think I have an explanation for that," or just, "Let me out of here"?

Diana: Well, there are two basic theories of the relationship between speech and music. One says that they're entirely modularized, so you've got a speech module that will only analyze the physical parameters of speech, and you've got a music module that will only analyze the physical parameters of music. Now, that obviously can't be true here, because the identical, physical parameters are heard in both speech and music.

Then, there's the opposite theory – that rather having been put into modules, what you hear is due to the physical parameters of the stimulus, whatever they happen to be. This illusion can't be explained by this theory either, since the physical parameters of the phrase don't change with repetition. So it seems that, whether we hear something as speech or as music depends, at least in part, on some top-down mechanism: something else is providing input into the system, and determining whether it should be analyzed as speech or music.

But in a sense both theories have some truth to them. There have to be modules involved at some level, otherwise we wouldn't be able to distinguish between speech and music. And on the other hand, speech and music must involve the same circuitry at some level – such as early in the auditory system.

Howard: Right, because otherwise you would never get confused.

Diana: Exactly. So, I think the bottom line here is that whether you hear something as speech or music is going to depend, at least to some extent – I don't want to exaggerate here because, obviously, there are some physical characteristics of speech and some physical characteristics of music—

Howard: Well, okay. I appreciate that, as an academic, you're worried about generalizing too much. But from my perspective, it seems to me that you're simply probing the boundary between the two.

We all know what's obviously speech and we all know what's obviously music; there's no debate there. But again, you've deliberately crafted this interference with the brain so that it seems to be imposing one structure when another one seems arguably more appropriate, and you're playing with these things to highlight the fact that there can be confusion within this processing system or processing systems, is that a fair way to look at it?

Diana: Yes, that's a fair way of saying it.

Howard: So then, once again, one starts to wonder, *What's going on?*

Diana: Yes. And you also start to wonder how often that might be happening in real life but we just don't notice it.

Howard: And what's interesting, at least to me, is that there are all these things that are taken for granted. One frequently hears people talk about "this system" and "that system", but what you're constantly doing, it seems to me, is saying, "Well, hang on, I can muddle these systems up; I can show you that these systems aren't as distinct as you might have thought. I can show you that we don't understand these things as well as we thought we might have understood them. They're much more complicated – who knows what the actual structure is and what's going on?"

And that, it seems to me, is the way science is done, by constantly probing things and highlighting problems. It gets back to what you were saying earlier about how some people feel uncomfortable about all of this. Now, on a personal level, I may feel uncomfortable because I might feel that I can't really trust my senses anymore; you're messing me up. But intellectually, it's incredibly stimulating because I'm thinking about how you're demonstrating all sorts of difficulties that I didn't even know existed. And, by so doing, you're pointing towards a deeper structure, or deeper level of our understanding, that must somehow exist.

Anyway, it's deeply fascinating. This is going to sound sycophantic, but I don't know why more people aren't doing this sort of thing.

Diana: I don't understand that either. I do stumble on these things. I couldn't have said, for example, "Okay, I'm going to invent myself an illusion called *Sometimes Behave So Strangely* now." It doesn't happen that way – instead I stumbled upon it. And I don't know why other people don't just stumble on things and go with what they hear.

I think it has to be what we were discussing earlier. I think that people don't necessarily trust their own judgment, or maybe think that there's something else wrong with them that they'd better keep quiet about. A lot of people don't like illusions, just like a lot of people don't like magic. I think it makes them feel uncomfortable; they don't want to think that they're being fooled. They like to feel that they're in control. That could be part of it.

Howard: Well, scientists are supposed to be better than that, but maybe they're not.

Diana: Well, they're interested in the work obviously; but I, too, am puzzled by why people don't just go ahead and create for themselves a lot more illusions than these.

Howard: Well, maybe they will. Maybe it just needs a little bit of time. Is there anything else that you'd like to add?

Diana: Well, we haven't talked about "Phantom Words" at all. Again, I was expecting to get an effect like the Octave Illusion but I didn't. It turned out that, if you have a two-syllable word or a two-word phrase and you alternate the syllables back and forth between the speakers, people start off hearing sounds that are word-like but aren't quite words, but after listening for a while, words and phrases just jump out at you.

Now, with that, I think we've pretty much covered everything.

Howard: Well, thank you very much Diana. This was a great opportunity to talk about some really fantastic work that you're doing.

Diana: Thank you very much. I've certainly enjoyed this a great deal, and learned a lot.

Further References

See Diana's webpage on the "Speech-to-Song" illusion.

See Diana's webpage on the "Phantom Words" illusion.

There is a wealth of material in the literature exploring the boundary between speech and music. Some examples include "Neural overlap in processing music and speech" by I. Peretz et al. (Phil Trans B; doi: 10.1098/rstb.2014.0090), "Predispositions and plasticity in music and speech learning: neural correlates and implications" by Robert Zatorre (*Science* Nov 2013:342(6158);585-589), "The relationship between music and language" by Lutz Jäncke (*Front Psychol.* 2012; 3:123), "Musical melody and speech intonation: singing a different tune" by Robert Zatorre and Shari Baum (*PLoS Biol.* 2012 Jul;10(7):e1001372), as well as a YouTube video of Duke University's 2014 Purves' 2014 talk, "Why We Hear What We Do: The Audition of Speech and Music".

Questions for Discussion

Introduction: Eclectic Beginnings

1. What do you think the general reaction of academics was to Max Mathews' music-generation software when he created it in the 1950s? Of musicians? Have things changed appreciably since then?

Chapter 1: Tones, Pitches and Critical Values

2. What were your expectations before you performed the tone-memory experiment with the intervening tones? Was it easier or more difficult than you anticipated?

3. What were your expectations before you performed the tone-memory experiment with the intervening numbers? Was it easier or more difficult than you anticipated? Did you feel any different, when performing this second experiment, compared to the first?

4. Does being told "not to attend to the intervening tones" have any effect, in your view?

Chapter 2: The Octave Illusion

5. What did you hear when you tried the Octave Illusion experiment? Did it remain constant?

6. Does this chapter make you regard other illusions, such as those from looking at an Escher lithograph, in a somewhat different way?

7. What does Diana mean when she refers to "top-down processing" towards the end of this chapter?

Chapter 2a: **Medical Applications?**

8. Do you think that future neuroimaging technologies will one day eliminate all invasive procedures like the Wada test? Is there a limit to what we might be able to learn about the brain from neuroimaging results alone?

Chapter 2b: **The Science of Eyes vs Ears**

9. Why was Diana motivated to conduct further experiments to "deliberately make things asymmetrical"?

10. What would be the advantage of studying stroke patients, as Norman Geschwind had suggested, to further probe the neurophysiological effects of the Octave Illusion?

11. What, if anything, do Diana's comments about the historical differences between studying hearing and vision imply about the difference in research orientation between psychologists and engineers?

Chapter 3: **Sociological Issues**

12. Do you agree with Diana when she says that most people suffer from a lack of self-confidence about what they hear, as opposed to what they see?

13. Are there many more musical illusions waiting to be discovered?

Chapter 3a: Auditory Fudging and Neuroplasticity

14. What did you hear when listening to the Scale Illusion? Do you find yourself gradually becoming more "comfortable" with these illusions now?

15. Are you surprised by the difference in processing power between the auditory and visual systems? Why do you think this is?

16. How might the increased likelihood of blind people having perfect pitch be related to the plasticity of the brain? Can it be otherwise explained?

17. Does our age of increasing specialization still allow for the possibility of hugely influential polymaths like Hermann von Helmholtz?

Chapter 3b: Surrounded by Illusions

18. Do you think that Tchaikovsky was aware of the illusion contained in his 6th Symphony? Was Nikisch? Might this sort of effect be happening more than we realize?

Chapter 4: Perfect Pitch and Tone Languages

19. Is it conceivable to you that people might have perfect pitch for speech, but not for music?

Chapter 5: Towards Monotony?

20. Is the gradual elimination of tone languages inevitable? Why might we think that this is not, in fact, the case?

21. Yuen Ren Chao's Chinese translation of Lewis Carroll's *Alice's*

Adventures in Wonderland is regarded as a classic, successfully transforming Carroll's notoriously clever English word-play into Chinese. Do you think it's possible, in principle, to fully translate *anything* from one language to another? Why or why not?

Chapter 6: Embracing Discomfort

22. Is there something special about the words, or phrasing, of "Sometimes behave so strangely"? Could you get the same sort of effect with almost any type of phrase? If not, what might the effect depend on, and what could that imply?

23. Which illusion was the most surprising to you? Which one was the most unnerving?

24. Do you believe that there is a deep, structural relationship in our brains between music and speech, or do you think that they are fundamentally different things?

25. Do you share Howard's confusion at why so relatively few people have gone into this field? Do you think this will change?

Topics for Further Investigation

Introduction: Eclectic Beginnings

1. How might studying theoretical models of the physiological basis of attention have affected Diana's later research orientation?

2. How has our understanding of attention changed since the "late-selection view" was first developed?

Chapter 1: Tones, Pitches and Critical Values

3. Diana mentions explicitly how technology affected the study of psychoacoustics. Can you give other examples where transformative technological advance shaped possible research avenues?

4. Why must this effect be somehow related to "lateral inhibition"?

5. How might this special two-thirds interval between tones be related to the "critical band"?

Chapter 2: The Octave Illusion

6. Do you imagine that something similar to the Octave Illusion would occur in everyday speech with a similarly constructed experiment? Why or why not?

7. Why might two separate pathways have evolved for the "*what*" and "*where*" processing of both sound and vision? What might be the evolutionary mechanisms responsible for such developments?

8. Do cube optical illusions, such as the Necker cube or Escher lithographs, act in a similar way (i.e. interfering with two different neural processing streams) to the Octave Illusion? Might we conceivably learn more about aspects of vision processing from auditory illusions?

9. Might it be a worthwhile exercise trying to combine auditory and visual illusions, or would that be entirely counterproductive?

Chapter 2a: **Medical Applications?**

10. What is our best current understanding of the relationship between handedness and cerebral dominance? What might be the mechanisms responsible to explain Diana's statement, "Handedness correlates very strongly with cerebral dominance, but not absolutely"?

11. Why is MEG "even less reliable than fMRI at present"? How might that change in the future?

12. If the Octave Illusion is a highly dependable probe of cerebral dominance, as Diana implies it might be, what are the relevant mechanisms at play? Might there be other such probes and how might they work?

Chapter 2b: **The Science of Eyes vs. Ears**

13. What might be the neuronal mechanisms involved in producing the Octave Illusion?

Chapter 3: Sociological Issues

14. Do you think that the study of musical illusions suffers from the prejudice of being somehow "unscientific"? If so, why might that be?

Chapter 3a: Auditory Fudging and Neuroplasticity

15. Might it be possible that illusions that trick our "unconscious inference", such as the Scale Illusion, still somehow depend on the distinction between two different auditory neural processing streams? How might the process of "unconscious inference" be concretely modeled?

16. How would you expect monkeys and humans to differ in terms of their auditory processing systems?

17. What, if anything, do you think that a rigorous understanding of auditory processing systems might imply about the Sapir-Whorf hypothesis?

Chapter 3b: Surrounded by Illusions

18. Why might it be easier to localize a tone that is richer in harmonics?

Chapter 4: Perfect Pitch and Tone Languages

19. Do you agree with Howard's suggestion that there would be an increased number of later-language learners of Mandarin with perfect pitch than would be found in the average non-tone language population? Do you agree with Diana that this effect would be

more pronounced depending on the age of learning Mandarin?

20. Are we all capable of developing perfect pitch? If not, why not? How might these sorts of investigations lead to a more precise quantification of the notion of neuroplasticity?

Chapter 5: **Towards Monotony?**

21. How might we conceivably test Diana's theory that music and speech evolved from a proto-language that included absolute pitch?

22. Would you expect that there would be a statistically significant difference between tone-language speakers and non-tone language speakers in their experiences of musical illusions? Which particular illusions would lead to more pronounced differences between the two?

23. Might the increased level of inference that tone-language speakers need to invoke somehow manifest itself in other non-linguistic ways too? How might we test that in principle?

Chapter 6: **Embracing Discomfort**

24. Do you agree with Diana that the Speech-to-Song Illusion provides evidence for "a very rapid and extremely long-lasting case of neuroplasticity"? If so, what might be the mechanisms responsible for that? If not, what might be an alternative explanation for the phenomenon?

25. How might the Speech-to-Song Illusion support Diana's above-mentioned thesis (see #21 above) that music and speech both evolved from a proto-language that included absolute pitch?

Ideas Roadshow
INTELLIGENT. INQUISITIVE. INTERNATIONAL.

OPEN AGENDA
PUBLISHING