[SQUEAKING] [RUSTLING] [CLICKING]

JOSH MCDERMOTT: All right. So today, we will continue our journey into the auditory system. Last time, we talked about the ear. And this time, we're going to be talking about one of the most important properties of the ear, namely frequency selective as well as the-- frequency selectivity, as well as the fact that the ear is nonlinear. And then we're going to get into hearing loss, which will probably affect all of you at some point in your life.

So just a refresher, this is the ear. Remember, sound enters here, travels down the ear canal, causes the eardrum to vibrate back and forth. Those vibrations are transmitted via the ossicles through to the cochlea, which is the organ that does the sensory transduction for hearing.

And so remember, we think of the ear as having these three main functional components. There's the outer ear, the pinna and the eardrum. There's the middle ear, which serves these two main functions-- impedance matching, which is necessary because you go from air to fluid, and overload protection. Remember, there are those muscles on the bones of the middle ear that can contract in the presence of high intensity sound. And then we spent most of our time last time talking about the inner ear. And we think of one of the main functions of the inner ear as being frequency analysis.

And here are some numbers that are worth knowing and fun to think about. So in each of your ears, there are approximately 3,500 inner hair cells. Remember, the inner hair cells are the cells that move in response to sound. And when they are moved, their membrane potential changes. So they transduce mechanical energy into a voltage.

There are about 12,000 outer hair cells. So remember that, typically, you have these three rows of outer hair cells, and then one parallel row of inner hair cells. So there's about a 3 to 1 ratio. And the outer hair cells are like the inner hair cells in reverse. So they get an electrical signal. And that causes their shape to change. Remember, we saw the video of the outer hair cell that was being supplied an electrical signal, which was music. And you saw the outer hair cell dancing.

The auditory nerve fibers are these essentially wires that are carrying information from the inner hair cells to the brain. And there's about 30,000 auditory nerve fibers per ear, so roughly about 10 per inner hair cell. So just for comparison-- and we'll talk a lot more about the visual system later in the class. But these are analogous numbers in each of your eyes.

So there are two types of photoreceptors in the eye. There are cones and rods. The cones, there are about 5 million of those. There's far more of the rods, 100 million. The rods typically function in conditions where there are low light levels. But the interesting discrepancy or difference with respect to the auditory system is that there are actually fewer optic nerve fibers than there are receptors, whereas, in the auditory system, it's the opposite.

So in the visual system, you're taking the super high bandwidth signal and squeezing it down into a narrow pipe. And that's what gets sent to your brain. And there's probably lots of consequences for the neural coding from that. And in the auditory system, it's just the opposite relationship. And we can talk in more detail later on about why that is. It's interesting to think about, but it's a salient difference.

So one of the most important attributes of the auditory system is frequency selectivity. So you saw this picture last time. These are tuning curves for auditory nerve fibers. So each one of these lines corresponds to a nerve fiber. And remember, each line plots the minimum sound intensity at a given frequency that is necessary to elicit a response in the fiber.

So you hook an electrode into a nerve fiber. You're playing frequencies of different intensities. And you're looking at what the threshold of the nerve fiber is. And so for a particular nerve fiber, say this one, there is a special frequency, which we normally call the characteristic frequency, to which it responds best. So the amplitude of the sound doesn't have to be very high in order to get a response. And if you change the frequency, higher or lower, the threshold goes way up.

So this empirical result suggests that the cochlea can be thought of as a bank of bandpass filters. And this property of frequency selectivity has a whole host of perceptual consequences. So this is just one. So how many of you are familiar with the phenomenon of beating? OK, a couple people. So if you play a string instrument you probably take advantage of beating when you tune your instrument. So beating is an acoustic phenomenon, and it happens anytime there are multiple frequencies that are present. And what's shown on the slide is an example of that.

So the slide shows you the waveforms of two pure tones. Remember, a pure tone is a tone that just consists of a single frequency. It's a sinusoid. So we've got the red pure tone that's at one frequency. And that sounds like this.

[PURE TONE]

Can everybody hear that? OK. And we've got the blue pure tone, which is a different frequency.

[PURE TONE]

It's a little bit higher. So you can see that, because those two pure tones are different frequencies, there's this point in time here where the phases align. So the peak of the red curve aligns with the peak of the blue curve. And over time, they drift out of phase. And so you get to this moment in time, where the peak of the red curve aligns with the trough of the blue curve.

Now, if you play those two pure tones at the same time, the waveforms add in the air. And so the sound that would actually occur if both of those frequencies were played at the same time is the black curve, which is just the sum of the red curve and the blue curve. And because of this particular relationship over time between the phases, you can see that the amplitude of the black curve waxes and wanes. So it's high here. And then by the time you get to this point in time, the red and the blue frequency, they cancel out.

So that's called beating, this waxing and waning in amplitude. And I mentioned that if you play a string instrument, you are probably familiar with that because you actually often listen to beating if you're trying to tune your instrument, because it'll tell you whether the notes that are played on two strings, if you play them at the same time, are exactly the same or a little bit different. So this is what it sounds like in this particular case.

[PURE TONES]

So it sounds a little bit like a telephone or something. So that modulation is perceptually the corollary to beating. And the word that we typically use to refer to that sound is roughness. So we call a sound rough when it contains amplitude modulation. So beating is an acoustic phenomenon. It happens any time you have two different frequencies that are being played at the same time.

However, you only hear the beating as roughness under particular conditions. And they turn out to be constrained by the cochlea. So the beats are only heard if the two frequency components fall within the bandwidth of the filters in your cochlea. And they also have to be in the same ear.

So if we take the red frequency-- if we use headphones, we take the red frequency and we apply that to your right ear and the blue frequency and we apply that to your left ear, you will not hear the beats. So you only hear them if they enter the same ear and if the frequencies are sufficiently close that they stimulate the same auditory filter. So here's just a quick demonstration of that.

So I'm going to play you three sounds. They each consist of two frequencies. But we're going to vary the difference between the frequencies. All right. So here, you have a one semitone difference. That's a 12th of an octave. So that's like the difference between a white and a black key on the piano.

[ONE SEMITONE]

And you can hear that. Here, it goes to three semitones.

[THREE SEMITONES]

And the beating becomes a lot more subtle. And here, it's out at eight.

[EIGHT SEMITONES]

And so at eight semitones, you can probably hear those two frequencies really as distinct things. You can almost attend to one or the other. And you don't hear the beating anymore. And it's not because the beating is not there. It's out there in the world. It's in the air. It just doesn't make it through your ear because those two frequencies get processed by different parts of the cochlea. So this is one respect in which the frequency selectivity of the cochlea really constrains what you hear.

OK. So one of the most important attributes of the auditory system is frequency selectivity. The tuning that we observe in the auditory nerve suggests that the cochlea can be thought of as a bank of bandpass filters. But the problem is that, in humans, we can't measure the filters directly.

So this picture that I just popped up, this is probably from a cat, or a rat, or a mouse, or some critter in which you can do an invasive experiment. And we don't have the ability to make these kinds of measurements directly in humans, because the auditory nerve is very, very difficult to get to.

So in humans, if you want to understand the frequency selectivity in the cochlea, you have to infer it indirectly from some other type of experiment. And the principal way that human frequency selectivity has been explored is by making use of what's called masking. So what is masking? Well, masking is defined as the process by which the threshold of audibility for one sound is raised by the presence of another masking sound.

So let's suppose there's a particular sound that you're supposed to be detecting. And you can hear it just fine. And then we play another sound at the same time. Under some conditions, that second sound will make it much harder to hear the first sound. And that's called masking. And this happens all the time in everyday life. And the reason for that is that sounds add in the air. And so any time there's more than one thing that's making sound, one of them is probably masking at least part of the other. So it's a very, very common phenomenon.

And it turns out that masking is very closely related to frequency selectivity. So this is a graph, an old graph, that shows the masking of a target pure tone-- so it's a little beep that you're supposed to detect-- by either a narrowband noise masker or another pure tone. And so what the graph is plotting is the amount of masking-- and I'll tell you in a second what that means-- as a function of the frequency of the target tone.

And there is a masker that is present at this particular frequency. And it can either be a noise band or a tone. And so the caption of the figure describes this in detail. So the pure tone is at 400 cycles per second, 400 hertz. And the narrow band of noise is centered at 410 hertz. So they're in this range here.

Now, the amount of masking, that's the amount by which you have to increase the level of the target tone in order to be able to hear it when the masker is also present. All right, what does the graph tell us? Well, it tells us, when the target tone has a frequency that's either much lower or much higher than the masker, well, then the amount of masking is very small.

You really don't have to increase the level of the target very much in order to be able to tell that it's there. But when the target is closer in frequency to the masker, you've got to crank up its level fairly significantly. And in particular, when it's very similar to the frequency of the masker, you might have to elevate its level by like 60 decibels. That's highly significant

So in general, the amount of masking is determined by the similarity of the frequency between the signals. Now, you might have noticed that the curve for the pure tone masker has got all these funny bumps and dips. And one of the reasons for this is that, when you have two tones that are present at the same time, you get this phenomenon of beating.

So if you're a participant in this experiment, trying to detect whether the tone is there, you can actually use the beating to help you detect the tone. So there are certain places where the tone gets a little bit easier to detect. So that's why you get these bumps and valleys in the graph, whereas, for the noise masker, you can't do that. And so you get a much smoother curve. But the key concept to take away from this is that masking is determined by frequency selectivity. And so we can use it in order to infer the shape of the filters that you have in your ears.

All right. And so the general model by which we do this is what's called the power spectrum model of masking. So this considers the cochlea to be a bank of filters. And you measure masking by asking people to detect a tone in the presence of a noise masker. And so the assumption is that a signal is heard-- so a signal here is the tone. It's heard when the signal to noise ratio at the output of the filter that you're trying to measure exceeds some value.

You're playing noise. And then you're playing the signal. And the idea is that some amount of the noise gets through the filter. And your ability to detect the tone will depend on the ratio between the power of the tone and the power of the noise heard through the filter that you're using to monitor the presence of the tone.

All right. So we assume that there is a filter that's characterized by a shape. The masker is a noise signal that will be characterized by some spectrum So this is power as a function of frequency. And the filter's filters got some response as a function of frequency. And so the power of the noise signal that gets through the filter is given by this integral. It's essentially the dot product between the spectrum of the noise and the spectrum of the filter.

And so the model posits that your ability to detect the signal will be a function of how much power gets through the filter. So let's just look at an example to build up some intuition here. So the idea is that there's this red thing in your ear. It's a filter. And we want to infer its properties. The noise signal is shown in gray. So this is frequency. So we're looking in the frequency domain. And this is plotting the masker power or the filter response as a function of frequency. So we've got a broadband noise. It's got all these frequencies in it.

So the amount of noise power that gets through the filter will clearly be influenced by the level of the masker. So if we just increase the intensity of the masker, the noise power will increase.

But it will also be dependent on the spectral properties of the masker. So here, we have a masker that's very narrowband. It's only got this narrow band of frequencies. And if we add low and high frequencies to make it more broadband, you can see that there's going to be more frequencies here that are hitting the filter. So the dot product of the filter and the noise will go up.

Now, the noise power will also be affected by the filter shape. So of course, we're assuming that there's a fixed filter shape in your ear. But we don't know what it is. And so one hypothesis might be that it would have this shape. Another hypothesis might be that it would have this shape. If the filter is wider, there will be more noise power coming through it. And thus, it should make it harder to detect the tone.

All right. So the general model that we're using here, in order to try to understand frequency selectivity via masking, makes a whole bunch of assumptions. It assumes the filter is linear. It assumes that, when you are performing this detection task of saying whether you hear the noise or not, you are only using one filter that's centered at the signal frequency.

So this is the signal. It's a pure tone. Remember, it's got one frequency. That's why there's a spike there. And in reality, your ear consists of this big bank of filters because the cochlea covers all frequencies. But we assume that you're just using one of them in order to detect the tone. And we also assume that the detection is based entirely on the overall power at that filter's output. And so none of these assumptions are strictly true. But they often provide what we think is a fairly reasonable first approximation.

All right. So the general method that would be used for measuring frequency selectivity in a person is to have some fixed signal, and then to manipulate the characteristics of the masker and to measure how the signal detection threshold will vary as the masker characteristics vary. And then we're going to try to deduce the characteristics of the auditory filter in your ear from the data.

So let me just give you an example I think will make this pretty clear. So this is like the classic experiment that was really the first demonstration of this approach. And it's attributed to Fletcher, who was a scientist who worked at Bell Labs back in the 1940s. And so people that were at phone companies were very interested in understanding frequency selectivity in the human ear and also masking, in part because masking really determines the extent to which you can compress an audio signal and have it still sound OK to a person.

And the essential concept is that, if you've got some sound with all these different frequencies in it, but you know that a whole bunch of the frequencies will be inaudible because of masking, you don't have to encode them. And so you can save bits in the signal that you're sending, say over a phone line. And so nowadays, all audio compression algorithms essentially have built into them measurements of human masking. So MP3 codecs have that built into them. So people were interested in this for a long time.

All right. So the experiment that Fletcher conducted involved measuring the signal detection threshold as a function of the noise bandwidth. So this is a noise that's got a constant spectral level. And we're just going to vary the bandwidth of the noise and measure how that affects your ability to detect a tone that's at the middle of the noise.

And so the results of the experiment are shown over here. So this graph plots the signal threshold in dB SPL as a function of the masking noise bandwidth in hertz. So over here, the masker is very narrow and the threshold is low. As you increase the bandwidth of the masker, the threshold goes up. And then there's this magical point at which, if you continue to increase the bandwidth of the masker, the signal threshold does not continue to increase.

And the intuition here is that, at this point, you've hit the bandwidth of the filter that you're using to detect the signal. So if you increase the bandwidth of the masker further, beyond the edges of the red curve, it's not going to affect your ability to detect the tone because those frequencies don't get passed through the filter. And that's called the critical band.

So we're going to demonstrate this and measure our own critical bands. What we're going to do is an experiment where I will play you a sequence of tones that will decrease in level. And your job is to count how many of them you hear. And the idea is the number that you can hear will be proportional to your-- inversely proportional to your threshold. All right, so let's check it out.

[AUDIO PLAYBACK]

- Critical bands by masking. You will hear a 2,000 hertz tone in 10 decreasing steps of 5 decibels. Count how many steps you can hear. Series are presented twice.

[DECREASING TONES]

[END PLAYBACK]

JOSH MCDERMOTT: Did you all get 10? Yeah? All right, good. OK. So you can hear 10. Now, we're going to do the same thing. But there's going to be a noise masker that will be present throughout. And so the tones will get harder to hear. And you won't be able to hear quite as many.

[AUDIO PLAYBACK]

- Finally, the bandwidth is reduced to only 10 Hertz.

[DECREASING TONES AND MASK]

[END PLAYBACK]

JOSH MCDERMOTT: I got eight. Is that about right? And I'm going to actually let you watch this. All right. So that was a very narrowband masker. But it elevated your threshold. So initially, you could hear 10. Now, after that, you could hear eight or nine. Is that consistent? Anybody disagree? No? All right. I hear some complaints from the crowd. All right, let's make it wider still.

[AUDIO PLAYBACK]

- Next, noise with a bandwidth of 250 hertz is used.

[DECREASING TONES AND MASK]

[END PLAYBACK]

JOSH MCDERMOTT: I got five. Yeah? Is that about right? OK. All right, so the threshold was elevated, went from eight or nine down to five. What's going to happen here, we're going to keep making it wider. And at some point, it's going to stop getting harder to hear the sounds.

[AUDIO PLAYBACK]

- Next, the noise has a bandwidth of 1,000 hertz.

[DECREASING TONES AND MASK]

[END PLAYBACK]

JOSH MCDERMOTT: I still got about five. All right? All right, so it's stable. Now, we're going to go to something that's a lot wider.

[AUDIO PLAYBACK]

- Now, the signal is masked with broadband noise.

[DECREASING TONES AND BROADBAND]

[END PLAYBACK]

JOSH MCDERMOTT: And what's kind of cool about that is that last noise is way louder than the one that came before it. But it's no harder to detect the tone. And the explanation for that is that you're using the filters in your ear to do that detection. And the response of the filter is not affected very much by all those extra frequencies that get added above and below it. So that's the critical band. All right. Any questions about this masking experiment or the general logic? Yeah?

STUDENT: Is it supposed to be that you hear more as it keeps getting wider? Because at least for me, I thought I heard more for the last one than the middle two. But I don't know if that's like an effect or not or just happens to be like that.

JOSH MCDERMOTT: My guess is it's a fluke. I mean, there's no real reason to think that that would happen. The fact that you actually hear a tone versus not is somewhat subject to random effects. Sometimes, you're paying a little bit more attention, or had a little bit more coffee, or whatever. And so from trial to trial, there'll be fluctuations. But yeah, on average, I don't think we would expect that you'd be any better there. It typically flattens out, yeah. Any other questions? OK.

The experiment here, it's conceptually very powerful. It makes a really good demo. And it's great for teaching purposes. But experimentally, in terms of actually using this to deduce the properties of the ear, it's actually not all that useful. And so nobody uses this anymore. And the main issue here is that the filter output power is really dominated by frequencies that are near the center of the filter.

And so imagine you're comparing this versus this. And because the filter has a weak response here in the skirts, the frequencies out there actually don't have a huge contribution to the overall amount of power. And so the technique is actually not super useful for providing an accurate estimate of the filter shape and bandwidth. So we mostly just now use this for teaching purposes.

So the standard modern method for getting estimates of frequency selectivity is what's called the notched-noise method. So the idea is what's shown here. So you're still trying to detect a signal. But now, instead of having noise that has this continuous spectrum, the noise has a notch in it. So you're supplying frequencies here and frequencies here. And then you can change the position of each of the bands of frequencies. So you can compare this, to say this, to say this.

And so you measure your threshold for these different types of maskers. And so the idea is that, in the black case, the only noise power that's coming through the filter is really the things that are out there in the skirts. And so you isolate the response at that particular point in the filter. And you can better characterize it.

All right. And so there's now been lots of experiments along these lines. I'm not going to work through in detail how you would go from the results of an experiment like this to inferring the exact filter shape. But you can work that out. And this is just an example of a derived filter function, where you have the response of this measured filter as a function of frequency.

And so the notched-noise method is generally acknowledged to be better than Fletcher's band-widening method. So it gives you more information about the filter skirts. It's also got this advantage that it reduces the potential for what's called off-frequency listening. So when you're doing this experiment, the whole point of the experiment is to try to characterize the red filter. It's the one that's centered at the signal frequency.

But the observer, the person who's doing the experiment, they have the option to use their entire ear. And in actuality, you have these other filters that are positioned at slightly different frequencies. And in principle, you might actually get better signal to noise ratios under some conditions listening to adjacent filters. But the idea of the notched-noise signal is that the filter that's actually centered at the signal frequency is the one where the signal to noise ratio would always be the best. And so the observer just doesn't have any reason to use anything other than that. So that's considered to be an advantage.

All right. So this has now been used fairly extensively in people to characterize frequency selectivity. So one of the main findings that comes out of these experiments is that frequency selectivity varies as a function of the center frequency. So this is a graph that is plotting absolute bandwidths. And that's the equivalent rectangular bandwidth. So it's one measure of how wide the filter is, measured in hertz, as a function of the characteristic frequency.

And so down here at 100 hertz, you can see that the bandwidth is like maybe 35 hertz. And then if you go up to a kilohertz, it's about 150 hertz. If you go up to 5 kilohertz, you can see that it's like 500 hertz. So as you move along the cochlea, the bandwidths vary. This is a pretty general property biological auditory systems, that the bandwidths tend to be higher at higher frequencies in absolute terms.

Now, if you look at the relative bandwidths, so if you express the bandwidth relative to the center frequency, they actually stay pretty constant or even decrease. And so these are filters plotted on a logarithmic frequency scale. And you can now see that they look pretty similar at different points along the spectrum and, in fact, maybe are a little bit narrower here at high frequencies than at low frequencies. And again, this is a fairly general property biological auditory systems. The filters have that property that, when you look at them on a logarithmic scale, they looked pretty constant.

So as I mentioned, there's lots of applications of these measurements of frequency selectivity. One of them has to do with using masking for audio compression. Another is that this is a pretty commonly used diagnostic to tell whether somebody's ears are healthy. So as we'll discuss, when you start to lose your hearing, the filters tend to get broader.

And then another application actually involves comparing the hearing of people to the hearing of non-human animals. And that's interesting for a whole bunch of reasons. Oftentimes, we're stuck with performing experiments on non-human animals because they're invasive methods that are not possible in humans. It's also really interesting, from the standpoint of thinking about evolution, to understand the relationship between humans and non-human animals.

And indeed, out of this type of work, there's emerged this interesting evidence that humans appear to have substantially better frequency selectivity than non-human animals. So it's like our ears are more tightly tuned. And so this is one example graph that shows an example of this.

So in this literature, they tend to not plot bandwidths. Instead, they plot the tuning sharpness, which is just reciprocally related to the bandwidth. So this is the tuning sharpness. And this is the characteristic frequency. And there are measurements that were made here in cats and guinea pigs. That's red and blue. And then the black lines plot two different estimates of human tuning bandwidth. And the point is that the black lines are substantially above the blue and the red lines. So the filters in the human are thought to be more sharp.

Here's another graph that shows cats, macaques, and humans. And here, you can see that the macaque monkey seems to be somewhere between cats and humans. And it's interesting to think about why this is. And really, there isn't widespread agreement on this. It's just an empirical observation. You might imagine that some of evolutionary pressure for communication might have driven this. Other people think it's just a function of the size of the cochlea. We don't really know. But it's an interesting difference.

All right. Now, I told you a little bit about some of these properties that we find in biological auditory systems. And it's natural to wonder why the ear is the way that it is. And some recent approaches that are now possible have given us some new ways of trying to understand these questions, where advances in machine learning now make it possible to optimize machine systems to solve interesting perceptual tasks.

And so this is a cool paper that came out of Google a few years ago now, where they, as the title says, learned the front-end of a neural network that had to recognize speech. I mean, Google was not really interested in the scientific questions that we're interested in this class. They were just trying to build a better speech recognition system. But the result of the study actually ends up being kind of informative with respect to thinking about the auditory system.

So this is a big typical system for recognizing speech. So you start out with a sound waveform. And in this case, the system convolves the waveform with a bunch of filters. And these filters, critically, were learnable. So there's convolution with these filters. There's some pretty standard operations of pooling and some nonlinearities. And then that's the input to some convolutional layers.

And then these things called LSTMs, long short-term memory-- the details of this aren't really critical. The point is it's a system that was optimized in order to recognize speech. And part of that optimization involved optimizing this front-end filter bank. And so you can do that. And then you can ask, well, what kind of filters emerge as a good engineering solution to the problem of recognizing speech?

So this is one way of looking at that. So these graphs are plotting the set of filters that are in a standard model of the cochlea that's based on biological observations, and then the filters that were learned by the system. So let's just focus on the model of the biological cochlea first.

So these are the filters ordered. And there's about 40 of them here. And you can see the frequency response of each filter. So this is frequency axis. And the gray scale here plots the response of that filter to the corresponding frequency. And the dashed green line is a function that's been fit to lots of observations from the masking experiments that we just talked about.

And so what you can see is that the filters in the ear are arranged from low to high frequencies with this particular mapping. But you can also see the thing that I just commented on earlier, which is that the high frequency filters are wider than the low frequency filters. So the point is that there's a wider black region here than here. So that's the same thing that we saw on this graph, where we saw that the bandwidth grew with center frequency.

And so the interesting result is that, if you look at the filters that emerged as a consequence of optimizing this system to recognize speech, if you arrange them in the same way-- I mean, it's a little bit messy, but you can see some traces of the same thing. So if you look at the filters that are tuned to low frequencies, they tend to be pretty narrow, whereas the ones that are tuned for high frequencies tend to be broader. And the distribution of frequencies is a little bit different. So you can compare the green curve here to the black curve here. It doesn't fall exactly on it. But it's not a million miles away either.

All right. So, so the inference that we might make from this is that the filters that are in our cochlea might be a good way to extract information from sound. Any questions about that?

STUDENT: What are the alternatives, like if you optimize it to do some different task, will it show a very different kind of--

JOSH MCDERMOTT: Yeah, that's a good question. This hasn't been looked at exhaustively. I can tell you that we've done a little bit of this in our lab with some tasks. And you can get-- if the task is narrow enough, you can get some unusual solutions that come out. And so I think it doesn't have to turn out this way. You could build a filter bank that wouldn't look like this at all, where you could build one that would just along the diagonal, where everything would be exactly the same bandwidth. That's just not what emerges from the optimization. Yeah. Any other questions? OK.

Now, we talked very briefly last time about how the tuning that you see in the ear changes with sound intensity. So in this respect, it's nonlinear. And this is another couple graphs that show the same of thing. So cochlear filter shapes, they change with sound level. And we'll talk in some detail about why that happens. But these are just two measurements that reflect this.

So this is physiological data. This is actually made from an animal cochlea, where you cut a tiny little hole in the cochlea. And then you can look in at the motion of the cochlea in response to sound. So you measure the motion of the cochlea as a function of frequency and intensity. And the results of the experiment here are expressed in terms of gain. So the gain is the ratio between the intensity of the stimulus and the mechanical motion of the basilar membrane. So it's a measure of how much the ear is amplifying sound.

And what this is showing is that, when the sound intensity is low-- so you can see the sound intensity is labeled here. That's what these numbers are. So that's 5 and 10 dB, 20, 30, 40. And so at low intensities, you get this very, very narrow tuning, where the gain is very high at a particular frequency range. This is frequency. And then it falls off. Now, as the intensity of the sound increases, the gain overall drops. But you can also see that the tuning gets broader. And so down by the time that you're at 80 decibels, the tuning is much, much broader than it was at 30 decibels.

What do these pictures show us? Well, they show us that the ear is providing amplification. And we believe that that's supplied by the outer hair cells. And it's level-dependent. So when your ear gets a very quiet sound, it's amplifying the heck out of it. When it gets a much louder sound, it's amplifying it much less.

And it's doing that in a way that is frequency-dependent. And so the consequence of that is that, at low intensities, the tuning tends to be pretty narrow. And at higher intensities, it gets broader. So this is a physiological measurement, again, made of the actual motion of the cochlea. This is psychophysical measurements made using the masking experiments that we've been discussing.

So each one of these curves is the measurement that's inferred from masking in a human with a particular signal frequency. So this is one signal frequency, this is another, this is another, this is another. And the different curves here are measured at different SPL levels of the signal. So this is dB SPL. So that's pretty quiet. And the highest one is 70 dB SPL. And you can see that the low intensities give you narrower tuning than the high intensities. So that's a psychophysical analog of that phenomenon.

All right. So in general, we see sharp tuning at low sound levels, broader tuning at high sound levels. And you can see this both physiologically and psychophysically. So another way to look at this level-dependent amplification is with a graph like this that shows the response of the cochlea, in particular the basilar membrane, as a function of sound pressure level. And when you look at it in this way, the basilar membrane exhibits what's called a compressive response function.

So what does this all mean? So we've got sound pressure here on the x-axis and two different measures of the basilar membrane response, displacement and velocity, which tell the same story on the y-axis. And they're both being plotted on log scales. This is in decibels. And this is a logarithmic scale.

And for reference, this line here displays the motion of the stapes. Remember, the stapes is the stirrup bone. It's the last bone in the ossicles that connects to the cochlea. So you can measure the motion of the stapes. And what you see there makes perfect sense. This is a line with a slope of 1. So that's a linear response. So if you increase the sound intensity by a certain amount, the motion of the stapes increases by the same amount.

Now, what happens on the basilar membrane is something that's very different. Instead, you increase the sound pressure by a certain amount. And the difference in the motion of the basilar membrane is substantially smaller. So it's a compressive response. Now, what's actually happening here is that there is amplification that is specific to very low intensity. So when you have a sound that's down here, 10 or 20 dB, the cochlea is amplifying the response to a pretty enormous extent.

And one way to think about this is to compare the response at the stapes to the response on the basilar membrane. So at high levels, you can see that the response at the stapes and the response of the basilar membrane are fairly comparable. But they diverge. So the point is that, at the stapes, the response to these very low intensities is like off the graph. It's like down here. But the response of the basilar membrane is way up here. That's because you have this amplifier in your ear that's boosting the response to these low intensities.

And it doesn't boost the high intensities nearly as much. And that's what gives you this compressive response function. So this is called compression. But really, it's level-dependent amplification. There's an amplifier in your ear. And it acts specifically when the sounds are low in intensity. And again, we believe this is the outer hair cells, which provide this kind of feedback system, whereby they get a little bit of input, and then they move, which, in turn, increases the motion of the basilar membrane. It's a little amplifier in your ear.

All right. Now, the story is a little bit more complicated than that. And that is that the amplification really only occurs for frequencies that are near the preferred frequency of a place on the cochlea. So these are, again, measurements of the basilar membrane motion. So it's the same kind of graph we were just looking at. But now, the different lines here-- and this is all made at one place on the cochlea. And the different lines correspond to different frequencies.

So this is made at the place that corresponds to 10 kilohertz, which is shown in the open symbols. And so at 10 kilohertz, you see this very shallow slope here for the cochlea's response. If you increase the frequency a little bit to 11 kilohertz, you still see a pretty shallow slope. But then as you go to 12 and to 13, the slope shifts to being pretty close to 1.

So as the frequency moves away from the characteristic frequency, the response becomes much more linear. So you stop getting this nonlinear amplification. So this picture is actually telling you really the same thing seeing as this picture. There are two different ways of looking at the same phenomenon.

So at low intensities, you're getting this massive amplification. But it's very frequency-selective. And so that's why, at low intensities, you get this very narrow tuning, because it's only the frequencies that the characteristic frequencies that get boosted. So at higher intensities, there's much less amplification happening. And as a consequence, the tuning gets much, much broader.

Any questions about this? You're going to have to go over this a few times. It just takes a little while to wrap your head around this. And so there's going to be a problem on the problem set where you think a little bit more about this. But the big picture here is that the ear is nonlinear. You've got this amplifier in your ear. And that is helping you hear these very low intensity sounds that we have to deal with on a daily basis. Yeah?

STUDENT: What is the characteristic frequency [INAUDIBLE]?

JOSH MCDERMOTT: So the characteristic frequency is the frequency to which a given place on the cochlea is most sensitive, responds the best.

STUDENT: [INAUDIBLE]

JOSH MCDERMOTT: In this graph or in this graph? It would be 10 kilohertz. So it's like you're getting the biggest response here at 10 kilohertz. So on this graph, it would be it would be 10 kilohertz. On this graph, each one of these fibers has a different characteristic frequency. It gives you the tip of the tuning curve. And it corresponds to the place along the cochlea. Yeah, go ahead.

STUDENT: [INAUDIBLE]

JOSH MCDERMOTT: Ah, that I don't-- you mean like with this converged to being right on the dashed line? I'm not sure. You might have-- there's probably some baseline level of amplification that you get throughout. Yeah, I'm not positive about that. But it gets pretty close, as you can see, yeah. Yeah, any other questions?

All right. Now, one of the consequences of the fact that you have these nonlinearities in your ear is that your ear creates frequencies that are not necessarily present in the input to the ear. And these are called distortion products. So the idea is that sinusoids are eigenfunctions of a linear system. So that means that, if you have a linear system, if you put in two frequencies, the output of the system will contain those same two frequencies, potentially scaled in amplitude and shifted in phase.

But if the system is nonlinear, then you can get additional frequencies that come out. And that's, in fact, what happens in the ear. So here's a situation where we've got f1 and f2. So we're supplying these two frequencies to the ear. And there are these other frequencies that we might potentially observe depending on the type of nonlinearity that's in the ear.

So if it's a quadratic nonlinearity, it involves a squaring. You'd get f1 minus f2. If it's a cubic nonlinearity, you'd get 2 f1 minus f2. And so you can think of any nonlinearity as being approximated by a Taylor series. So there'd be a quadratic term and a cubic term. And those tend to be pretty prominent in our ears.

So we would expect that, if you've got this like powerful nonlinearity in your ear, it's going to cause these new frequencies to be present. How would we potentially detect these? So this is a classic demonstration. That's one of the main ways in which we actually measure distortion products. And it's going to make use of beating, which we just talked about at the start of the lecture.

All right. And we'll listen to this demonstration. But I'm going to explain it first. So we're going to play you a sound that contains two pure tones, one at 1,000 hertz and one at 1,200 hertz. So the cubic distortion product, which we would predict would be present in your ear if there is a cubic component to the nonlinear function in your ear-- the cubic distortion product would be 2 times f1 minus f2.

So 2 times f1 would be 2,000 minus f2 is 800 hertz. So the theory predicts there should be this additional frequency that's being created in your ear by its nonlinear properties at 800 hertz. So to probe for the presence of that, we're going to add a pure tone at 804 hertz.

Now, the logic here is that, if your ear is creating 800 hertz and we add an 804 hertz, there's going to be beating. Now, what I didn't tell you before is that beating always happens at the difference of two frequencies. So if we've got 800 hertz and 804 hertz, there should be beating at 4 hertz, four times a second. So you'll hear this low frequency fluctuating wobble. And so we're going to actually do this. And you can see whether you hear this. And I'm going to show you this, so you can see what you're listening to.

[AUDIO PLAYBACK]

- Aural combination tones-- in this demonstration, two tones of 1,000 and 1,200 hertz are presented. When an 804 hertz probe tone is added, it beats with the 800 hertz aural combination tone.

[AURAL COMBINATION TONES]

[END PLAYBACK]

JOSH MCDERMOTT: All right, so the point is that, when you just have the two primaries present at 1,000 and 1,200, you shouldn't hear this subtle fluctuation. Then we add in the third tone at 804 hertz and you should hear this subtle pulsation. So that's how you're supposed to tell whether you've got the distortion product or not. So we can just try it again.

[AUDIO PLAYBACK]

[AURAL COMBINATION TONES]

[END PLAYBACK]

JOSH MCDERMOTT: All right. So it's pretty subtle, but it's there. And in a second, I'll play you something that's not as subtle. So this is actually one of the ways that people actually measure these things. And there's something even crazier that you do, if you actually want to infer the amplitude of the distortion product.

And that is-- so you do what we just did, all right? So you've got your two primaries, and then the beating tone. And then you add in a tone at 800 hertz. And the participant has to adjust the amplitude and phase of the 800 hertz tone that you're piping into their ear. And the idea is that, if they can get it exactly out of phase with the frequency that your ear is creating, it cancels out, destructive interference. And then the beating goes away.

And so you do these incredibly painful experiments, where you're twiddling these knobs, adjusting the amplitude and phase of this 800 hertz thing to make the beating go away. And so we did this once in our lab. And it's a paper that has only 1 and 1/2 participants. One of the participants was the grad student who led the study. And he did it in both of his ears. And then the other participant was another student. And he could only bear to do one of his ears because it's really, really draining to do. But you can do it. And you can measure these things.

But we have an even cooler demonstration of this, which is that we're going to-- let me turn down the volume here. All right. So we expect that there is a distortion product that's 2 f1 minus f2, the cubic distortion product. And so the idea is that, because there's a minus sign in front of the f2, if we actually induce a frequency sweep-- so we cause f2 to sweep up-- the distortion product should sweep down in frequency.

And so that actually makes it really easy to hear because you can tell that there's this thing that you're hearing that's going down when the other stuff is either staying stationary or going up. So this is what you'll hear.

[F1 TONE]

So that's f1.

[F2 TONE]

That's f2. And now, the both.

[F1 AND F2 TONES]

You hear it going down? OK. OK, let's do it again.

[F1 AND F2 TONES]

So you can see how the stimulus is going up. That thing that was going down is in your ear. It's not in the signal. So that's a distortion product that your ear is creating. Any questions about distortion products? Yeah?

STUDENT: I don't understand what it means to have something created in the ear. Is it like a perceived thing? Are we creating it not in our minds? Like, how does it have physical tangibility in the ear?

JOSH MCDERMOTT: Well, yeah, I use the word created because the nonlinearity in your ear, it actually is reflected in the motion of the basilar membrane. So your basilar membrane will actually be vibrating at the frequency of the distortion product.

STUDENT: [INAUDIBLE]

JOSH MCDERMOTT: That's correct, yeah. Yeah. Yeah, any other questions about that? OK. So take-home messages are the ear is frequency selective. You can't measure it directly in humans. You have to infer it from masking. So we talked a little bit about how you do those masking experiments. The results of the masking experiments show that frequency selectivity varies with frequency.

So at high frequencies, it tends to be broader than at low frequencies. We talked about some evidence that maybe that's a good way to design an ear in order to help you perform auditory tasks. We also talked about how frequency selectivity changes with level. And that's because there is amplification that's happening in your ears that is both level-dependent and frequency-specific. So it happens at low sound levels at frequencies that are close to the characteristic frequency. And the consequence of that is that your ear is a nonlinear system. And that is evident in these distortion products that you can hear if you do clever things.

And so I should also emphasize about these distortion products-- I mean, these are there all the time. It's just that most of the time you don't notice them. It's just part of the way the world sounds to you. And you have to do these funny experiments in order to see very, very clear evidence of them.

OK, let's start talking about hearing loss. So this is a graph that shows normal auditory sensitivity. So this is plotting detection thresholds expressed as the sound intensity that you can just barely detect as a function of frequency. So it has this U shape. So you're most sensitive typically around 2 or 3 kilohertz. And then when frequencies get lower, you become less sensitive. And when they become higher, you're less sensitive. So this is just what somebody with normal hearing looks like. This is what gets measured when you go to the doctor. You want to pass those out? Yeah, sure.

STUDENT: [INAUDIBLE]

JOSH MCDERMOTT: Ah, we have some hard copies of slides. There's a few people that wanted those, yeah. OK, so this is what gets measured. If you go to the audiologist and you have your hearing checked. We do this in our lab. This is me getting my audiogram measured this morning thanks to a grad student who helped me do this.

And so this is what it looks like if you go to the audiologist. These are frequencies here. And this is the threshold here. And I mentioned in the previous lecture that we often talk about hearing level. So there's often SPL measurements that are absolute level measurements relative to the standard.

And then hearing level is expressed relative to what's considered normal hearing. So 0 would mean that you have perfect hearing. And my hearing is not perfect. But it's not terrible. It could be worse. And you can see here that, at higher frequencies, I've got some significant threshold elevation, especially in the left ear. So at 4,000 hertz, my threshold is like 35 dB higher than presumably it was when I was born.

So if you go to the doctor, this is the kind of thing that they'll do. If you're a hearing scientist, you often do this before you run somebody in an experiment, just to make sure their hearing is normal. And this is what they look like as a function of age on average. So this is that same thing, but now averaged within age groups.

So we've got the threshold expressed in dB HL. So again, 0 is considered normal. And you can see that, in your 20s, which is the black curve and 30s, the red curve, life is pretty good. On average, there's very, very little threshold elevation. But as you move beyond that-- 40, 50, 60-- there starts to, on average, be some fairly significant threshold elevation. And it's just systematic with age. And it's really quite significant. So you can see that people in their 80s, like at these high frequencies, 4 and 8 kilohertz, we're talking decibels of threshold elevation, so pretty substantial. So this is a massive public health problem because it basically happens to almost everybody. And there's lots and lots of interest in trying to understand both why this occurs, how to treat it, whether to prevent it, things like that. Yeah?

STUDENT: So does this graph-- as you go higher, does it gradually go up? Or does it asymptote at some point once you get like, higher [INAUDIBLE}. If you extended this frequency all the way to like 20 kilohertz does is it-- is it an asymptote?

JOSH MCDERMOTT: No. So it would--

STUDENT: [INAUDIBLE]

JOSH MCDERMOTT: Yeah, it'll look--

STUDENT: [INAUDIBLE]

JOSH MCDERMOTT: Oh, well, you just-- that would be hard to measure. But yeah, no, you wouldn't really bother measuring anything that's more than 80 decibels because it's just not really practically very relevant. But yeah, I mean, it'll look worse. So I mean, typically if you take somebody in their 30s, they will have threshold elevation in the high teens of kilohertz. Yeah, the very high frequencies are the most vulnerable. And that's where you see the biggest threshold shifts. Did you have a question?

STUDENT: Yeah, I was going to ask if there's any hearing loss that's characterized with losing lower frequency sounds, and then still being able to hear the higher frequencies.

JOSH MCDERMOTT: Yes. It's more unusual and typically associated with particular genetics. But yeah, you do you do find people like that. Yeah?

STUDENT: What's the deal with the 40 and 50-year-olds?

JOSH MCDERMOTT: I think that's a midlife crisis. No. I mean, I think the sample size here is just-- it's not infinite. And so there are error bars on these. But yeah, we would certainly expect that this would be monotonic with age if you had enough data, yeah. Yeah?

STUDENT: [INAUDIBLE]

JOSH MCDERMOTT: So we'll talk about conductive losses in a second. Those are not so common. So this is really just reflecting typical age-related hearing loss that just happens to almost everybody. There are specific diseases that would have different phenotypes that are less common. Yeah?

STUDENT: Is this true for people living in all places or if you're rural--

JOSH MCDERMOTT: Stay tuned. Yeah, great. That's a great question, yeah. OK. So just to give you a sense of the significance of this, this is the amount of attenuation that you would get from having earplugs in your ears or just putting your hands over your ears. So this is sound frequency and the amount of attenuation. So earplugs maybe give you 20 or 30 dB attenuation more at high frequencies than at low frequencies. Putting your hands over your ears gives you that. So this is on par with or maybe not even as extensive as what you would have if you just get older on average.

So there's two kinds of hearing loss that we will typically distinguish between. I hope you don't do that the whole lecture. So there's conduction loss, which we were just talking about. So that's an impairment in the mechanical transmission of sound energy to the inner ear. So it's typically a problem with the outer or middle ear. So for instance, if the bones in your ear don't move as well as they should, you won't get as much energy into the cochlea as you should. And your hearing will be worse.

So that's contrasted with sensory/neural loss, which is an impairment in the transduction of sound energy to electrical signals or in their transmission to the brain. And that's usually due to problems with the inner ear. So conduction loss can involve the puncturing of the eardrum-- for instance, via an ear infection.

There's also otosclerosis, which was just asked about. So that's a disease that involves the gradual immobilization of the stapes bone and the ossicles. So it just doesn't move as well. It's inherited. And the good news is that this can usually be treated pretty effectively. So they can perform surgeries that will replace the stapes bone. And that's a pretty common surgery for an ear surgeon to perform. Yeah?

STUDENT: [INAUDIBLE]

JOSH MCDERMOTT: I forget. I'm not sure. Yeah, I can find out. So sensory/neural loss is the more vexing problem that most of us deal with. So one aspect of this is what's normally referred to as presbycusis. So that's the thing that was just on that graph that we were showing, which refers to the fact that, as people age, they become less sensitive and particularly less sensitive to high frequencies.

So we didn't see this on the graph. But as I was alluding to, most people over 30 typically can't hear above 15 kilohertz. Most people over 50 can't hear above 12 kilohertz. So did any of you-- like when maybe you were in high school, did you have those high frequency ringtones that your parents couldn't hear? Or you've heard of them? Your friends have them? Anyhow, this is a thing, right? You can program your phone to have a very high frequency ringtone, if you don't want your parents to know that you're getting phone calls. And it's because older people can't hear these really high frequencies.

OK, so it's unclear what causes presbycusis. There appears to be a genetic component. And part of what makes this-- part of what makes this hard to understand is that while people are alive, it's really hard to actually see inside their ears. So you can measure the fact that somebody's hearing is worse, but you can't really look inside their ear, their auditory nerve, and actually directly observe what would be different about somebody who has poor hearing compared to somebody who has great hearing.

So one of the main ways that this has been studied is what is called otopathology. This is postmortem studies of the cochlea. So when people die, if they donate their bodies to science, the cochlea can be removed. You can chop it up, put it under a microphone, and then look at it in all kinds of different ways.

And so one of the classic things that has been done is to actually count the number of hair cells that are present in a person when they die. So this is an example of a slice of the cochlea. This is the organ of Corti, so basilar membrane here. And we've got outer hair cells here and an inner hair cell here.

So you look at these things under a microphone and you count them up. And so these would be different slices along the cochlear spiral. And you just count the presence or absence of outer hair cells. So you see three here. And inner hair cells, you see one there. So present, absent as a function of distance from the base.

In fact, our neighbors across the river at the Mount Sinai Ear Institute are a worldwide leader in this kind of thing. So they have lots and lots of data on this. And these slides are actually borrowed from my colleague, Charlie Lieberman, who works there.

So these are results of a couple studies that did this. So they get these cadaver cochlea. And they know the age at which the person died. And it can be just anywhere on this range. So each dot here is a person. And the graph here is plotting the percent of hair cells that are present in that individual at death.

And we have IHC-- that's inner hair cells-- on the left and OHCs-- that's outer hair cells-- on the right. And the different colors are just two different studies, one in 1968 and one from a couple of years ago. And what these graphs show is that, as the age of an individual increases, at death, they have fewer hair cells. So there's a few people here who, unfortunately, died very young. And you can see that they have almost all of the hair cells remaining. And then it, on average, gradually drops off with age.

So this is a graph that plots the hair cell survival as a function of position on the cochlea. So the previous graph is just pulling everything together. And here, we have hair cell survival as a function of position on the cochlea. And what this shows is that, for inner hair cells, there's considerably more hair cell loss at the base of the cochlea.

And which frequencies get transduced at the base? High or low? High, that's right. So base transduces high frequencies. So that's consistent, in principle, with the fact that people tend to lose their hearing at high frequencies more than low. You can see that outer hair cells also tend to be lost, but it's more distributed across the cochlea. So both the base and the apex have outer hair cell loss and even in the middle.

So this is a study. So you can also look at this in lab animals. And these are graphs that plot five different species of lab animal-- gerbils, chinchillas, mice, rats, and guinea pigs. And this is their outer hair cell survival as a function of position along the cochlea. And if you compare that to humans, there's this interesting difference, that humans seem to have considerably more outer hair cell loss at the base of the cochlea.

Now, what is it that's different about humans than all these lab animals? Well, the lab animals lead very quiet lives. They sit in a boring room in a lab and are not exposed to all of the noise and rock concerts and stuff that normal humans get exposed to. So you might imagine that this is consistent with the idea that some component of this is related to noise exposure.

So there's this hypothesis that prolonged exposure to noise, possibly even that which occurs Incidentally in industrialized societies, likely plays a role. And so this gets back to the question that was asked a little bit earlier, which is what happens in other places. So there's a couple of intriguing studies that have been done that suggest that, in non-industrialized places, there's quite a bit less of age-related hearing loss.

So this is one study that was done in 1962 in a fairly remote area in Africa. So it says, in the description here, the area chosen is about 650 miles southeast of Khartoum, the capital of the Republic of the Sudan, a few miles from the Ethiopian border and 10 degrees above the equator. Until 1956, this area was a closed one, untouched by any foreign culture or civilization.

So they went to this place and measured audiograms. And the results are shown here. So this is this group of people called the Mabaans. And they compared the results to a very large study of hearing that was conducted at the Wisconsin State Fair in 1954. They set up a booth at the Wisconsin State Fair. And all these people came and had their hearing measured.

And so people in Wisconsin in the '50s show something that's consistent with the graph that I showed you earlier. So as a function of age, hearing gets worse, especially at high frequencies. So the Mabaans, on the other hand, show much less of this age-related hearing loss. Yeah?

STUDENT: What does it mean when [INAUDIBLE] negative?

JOSH MCDERMOTT: Yeah. I mean, the 0 is just-- it's like a standard. It's supposed to be like the average person who has good hearing, normal hearing. So you can do better than that, yeah. So you might think this is circumstantial evidence that, well, there's presumably some difference between the people in Wisconsin and the Mabaans that accounts for this. And there could be lots of things. But one candidate is the fact that industrialized society is very noisy. We're around noisy machinery all the time. And that might not be good for our hearing.

So the other study that was done along these lines was conducted on Easter Island. This is an island that's off the coast of South America. It has been called the most isolated inhabited place in the world. And I guess, at some point-- I don't know if this is still the case. But it was known as the Island of the Great Silence because it's non-industrialized. So you don't have modern day machinery there. At least, you didn't at the time when this was done. And so what they did is they went to Easter Island and measured audiograms of people there.

Now, people who live on Easter Island, I guess, would occasionally go to mainland South America for some period of time, like to work for a little bit, and then sometimes come back. And so the researchers divided the participants in this study into those that had never left Easter Island, and to those that had been away on mainland South America for three to five years, and those that had been away more than five years. And so this is the average audiogram for those three groups of people.

So what this shows is that the people that had been on mainland South America for some period of time had worse hearing than the ones that had never left. Interestingly, if you take the group of people that had never left and you subdivide them according to age, you still see signs of age-related hearing loss. So the older participants have worse hearing than the younger participants. But this, again, is circumstantial evidence that something about living on mainland South America causes your hearing to deteriorate. And it could be noise exposure.

So there's this hypothesis that the incidental noise exposure that happens with modern life may contribute to age-related hearing loss. And I think it's plausible that we evolved in conditions that were just much less noisy. And so the ears are just ill-equipped for the sound levels that we encounter in life today. That seems possible.

Now, what's much less controversial is that acute noise exposure can cause hearing loss through damage to the organ of Corti. So if you get exposed to something that's really, really loud-- it could be gunshots, could be an airplane, if you're an airplane mechanic-- we have pretty clear evidence that that's not good for your hearing.

And some of the clearest evidence actually comes from this really interesting study of hunters. So this is a person holding a rifle. And in this particular study, they recruited people who, I guess, were aspiring hunters. And they measured their audiograms before they learned how to hunt, and then after they had been hunting for a while.

So the person's holding the rifle-- and the idea is that there's some-- you fire the gun. And there's an explosion in here that causes a big sound wave to come out. And there's a direct path from the source of the sound wave and the left ear, whereas the right ear is shielded from the sound.

And what you can see here are the pre-exposure audiograms. Those are the solid lines. So the x's are the left ears and the o's are the right ears. And then the post-exposure audiograms-- so you can see that post-exposure-- so that means after the person started hunting-- the thresholds are elevated. But they're substantially more elevated on the left than the right, consistent with what you would expect from the noise exposure.

OK. Eventually, the people who became hunters passed away and donated their ears to science. And this is the result of otopathology done on their ears. So we've got the left ear and the right ear. And these are the counts of inner hair cells and outer hair cells. And so there's three lines, because there's three rows of outer hair cells typically, along the length of the basilar membrane.

And you can see that, at the basal end of the basilar membrane in the left ear, there's this precipitous drop in the presence of hair cells. And certainly, the right ear also doesn't look normal. But it's not nearly as pronounced as what you see in the left ear. And we, again, think that's because the right ear is protected by the acoustic shadow. Any questions about that? Yeah?

STUDENT: So have there been studies where lab animals have been exposed to noise and their hearing is measured?

JOSH MCDERMOTT: Yeah, there's lots of those. In fact, that is what we are going to talk about next. Yeah, so we know quite a lot about the temporary or transient effects of noise exposure from these studies of lab animals, yeah. Any other questions about this? So OK, before we wrap, if there's one thing that you take away from this class, wear earplugs. You're going out to hear music? Wear earplugs. You'll be happy when you're older.

And in fact, the thing that I'm going to tell you next, which will have to wait until next time, is going to be some particularly sobering evidence that actually hearing protection is important, because one of the other things that can happen is, in addition to these threshold shifts that show up in the audiogram, so where your thresholds get elevated, there are also temporary threshold shifts.

So if you get exposed to moderately loud sounds-- this is like the rock concert effect. You go out to a nightclub or a concert. You probably have this experience that you leave and you can tell that your hearing is not normal. It's harder to hear. Maybe you have some ringing. And that's measurable.

So these are temporary threshold shifts. So if you went and measured your audiogram right after you went to the nightclub, the thresholds would be elevated. But then they go back to normal. You wait a day or even a few hours and they go back to normal. But as we will see next time, there's reason to believe that, even though your thresholds return to normal, your auditory system has potentially been altered by this. And we will hear about that next time.