

Chapter 4

Human Echolocation

(with Michael S. Gordon)

[B]at sonar, though clearly a form of perception, is not similar in its operation to any sense that we possess, and there is no reason to suppose that it is subjectively like anything we can experience or imagine.

– Thomas Nagel, “What Is It Like to Be a Bat?” (1974, p. 438).¹

i.

Hold this book open before you and read this sentence aloud. Can you *hear* where the book is? Can you hear that it has a certain size, shape, and texture? Pull the book away and continue to speak. Can you hear the emptiness of the space before you? If you close your eyes and speak, can you hear that there’s a wall a few feet to your left, a large desk at hand level?

Michael Gordon (co-author of this chapter and hereafter “Mike”), a psychology professor who studies sensation and perception, has convinced me that I can and often do hear such things – that I’m more bat-like in this way than I had previously supposed, more bat-like than Nagel, in the epigraph, takes us to be. We suspect that you, the reader, if you’re not substantially impaired, are also somewhat bat-like in this way.

Unless you’re already acquainted with the relevant literature, you probably doubt that you can hear silent objects and their properties, or at least that you can hear them

very much. By the end of this chapter, Mike and I hope that you'll come to regard such doubts as a naive mistake – a twofold mistake, in fact – a mistake not only about your sensory capacities, that is, about your ability to respond to inputs in a certain way, but also about your stream of conscious experience, your everyday auditory phenomenology, including the phenomenology you (probably) think you had or didn't have when you tried the experiment described in the first sentences of this chapter. Mike and I will argue that people typically know only poorly the auditory phenomenology produced by silent objects, thus providing another case in support of the central contention of this book, that people are, in general, poor judges of their own stream of experience.²

ii.

First, let's look at the basic experimental research on human echolocation. As Mike and I will be using the term, echolocation is the ability to detect features of the environment, especially features of objects that generally do not themselves produce sound, using the acoustic changes in sounds from other sources as they reflect off or are otherwise mediated by those environmental features or objects.³ When a bat detects the presence of a (silent) wall by gauging how sounds reflect from it, it is echolocating. Likewise, if a bat can learn about a bug using information about how sound is transformed in passing through the bug, that too is echolocation. In echolocation, one detects properties of objects not by detecting waves that emanate directly from those objects as sound *sources*, but rather by detecting how those objects act as sound *reflectors* and sound *modifiers*.

Many species of bat, of course, and dolphin, and whale, echolocate very accurately (classic treatments include Griffin 1958; Kellogg 1958; Evans 1973).⁴ Blind people also often use echolocation while they walk – or even bicycle! – through unfamiliar or changing environments, often tapping a cane or making clicking sounds with their mouths (Supa et al. 1944; Cotzin and Dallenbach 1950; McCarty and Worchel 1954; Rice 1967; Stoffregen and Pittinger 1995). The blind mobility instructor Daniel Kish (2009) offers a delightful retrospective account of life as blind child, using echolocation to navigate by foot and bicycle, to play tag and climb trees. Several videos demonstrating the tremendous echolocatory abilities of the blind, including bicycling and rollerskating, can be found at www.worldaccessfortheblind.org.

A small body of empirical research shows that people with normal vision can also echolocate, at least a little bit, with brief training. For example, Michael Supa and colleagues (1944) asked both blind and normally sighted people to walk toward a large masonite board mounted at various distances from them. During their approach, participants signaled both the moment when they first detected the board and also when they were as close to the board as possible without touching it. Blind participants could detect the board several feet before contact and could move to within a few inches of its surface. After about thirty trials, sighted participants achieved similar accuracy. To confirm the auditory basis of this ability, Supa and colleagues reduced auditory input – first having participants remove their shoes and walk toward the board in socks, then giving them earstops, then finally projecting noise directly into their ears. People’s performance deteriorated proportionately with their induced deafness to the point where

all participants collided with the board in every trial. (They were walking slowly of course, so no harm was done.)

Carol Ammons and colleagues (1953) and Lawrence Rosenblum and colleagues (2000) found similar results. Sighted but thoroughly blindfolded participants, after brief training, were able to stop the moment before walking into a large, sound-reflecting surface. Rosenblum and colleagues also found that sighted participants could use echolocation to discern the approximate distance of a wall that was positioned from three to twelve feet in front of them. Blindfolded participants echolocated using self-generated sounds (such as saying “hello” repeatedly), and possibly ambient sound as well, while either moving or standing still. Then, with the wall removed, they estimated its distance by walking to where they thought it had been. Although the task was easier while moving, even stationary participants had some ability to detect the three foot differences between wall positions.

Steven Hausfeld and colleagues (1982) asked blindfolded, sighted participants to echolocate an object placed 25 centimeters before the face. They varied the texture and shape of the targets (fabric, plexiglass, carpet, or wood; circle, triangle, or square [all of equal surface area], or no target). After brief training, participants could distinguish – not dependably, but better than chance – between the shapes and between some of the textures (see also Rice 1967; we also informally replicated this in my office, finding some people to be approximately at chance and others to be more than 50% accurate in shape detection). Rosenblum and Ryan Robart (2007) found that blindfolded participants could distinguish, at rates above chance, between triangles, squares, and circles, using noise emitted from speakers positioned directly behind the shapes. Gordon (that is,

Mike) and Rosenblum (2004) found that people could judge the size of an aperture – specifically, whether they’d be able to walk through it without turning their shoulders or ducking their heads – while blindfolded and hearing crowd noise from speakers behind the aperture.

How are people doing this, exactly? The question has only begun to be studied. The most deflationary interpretation, perhaps, is that people are simply noticing differences in sound intensity: Their “hello” sounds louder when reflected back from a nearby wall than when uttered into empty space; a smaller aperture permits less crowd noise to pass through. This is probably too simple an interpretation, though.⁵

Daniel Ashmead and Robert Wall (1999) suggest that the most important cue for avoiding walls may be the accumulation of low frequency sound, in particular. Other possible sources of information include the time delay between an emitted (or otherwise localizable) sound and the return of its reflection; differences in loudness, pitch, and timbre due to patterns in the efficiency of the reflection and transmission of different acoustic frequencies; and interference patterns in reflected sound.

This body of research demonstrates beyond a reasonable doubt, we think, that ordinary, sighted people at least *can* echolocate, in certain conditions, with a bit of training. People can detect the presence or absence, and to some extent the distance, maybe even the shape and texture, of silent objects by hearing how sound reflects off, transmits through, or reverberates from those objects. Of course, whether people actually *do* echolocate in their daily lives, and whether there’s an auditory experience of echolocation – those are separate questions, to which we now turn. We’re unaware of any systematic research on these broader questions; so we’ll have to employ

introspection, anecdote, and plausibility arguments. In other words, it's time for a bit of fun. (Okay, we admit to having a slightly perverse sense of fun.)

iii.

First, consider your experience walking down a long, tiled hallway in hard-soled shoes. With each step, a burst of noise radiates into the area around you. In hearing this sound, you hear not only the shoe striking the floor but also the reflections of that sound from surrounding surfaces. If the space were much different – if, say, you were taking a few steps across the tile of a bathroom floor – your auditory experience would be quite different. Similarly, there's an obvious echoic difference that makes you sound like Pavarotti in the shower and like yourself everywhere else. Hallways and showers sound different. They do so not because hallways and showers *produce* different sounds but because they *reflect* sound differently. In reacting to their acoustic differences, you're using echoic information.

If what we think to be a concert hall doesn't sound like a concert hall or what we think to be a shower doesn't sound like a shower, we'll ordinarily notice the difference – or so, at least, it seems plausible to suppose. If you were in the shower embellishing a cadenza and suddenly, silently, the walls were removed so that your echoic environment became that of a concert hall, we expect you'd be rather startled. More subtly, imagine stepping through a doorway into a familiar tile hallway, visually focusing to the left, and being surprised by the sound of your footstep. Turning right, you discover a large piece of furniture where none was before. Echoic information would, in such a case, be

guiding visual attention. If such examples are telling, then in moving through the world, we're constantly using echoic information at least to supplement and confirm what we know primarily through sight, memory, and non-echoic aspects of hearing.

Wenger Corporation has developed what they call a "virtual room". This room is able to artificially synthesize the acoustics of a variety of spaces, from an office to a symphony hall (something that practicing musicians have found useful). If the acoustics of the virtual room are set to emulate an area much larger than the actual size of the room, listeners will quickly notice that something is amiss. Typically, people entering the room will glance upward to see if the ceiling might be especially high. If echoic information weren't regularly being used to supplement other sensory information, people wouldn't react this way, since of course people don't normally glance at the ceiling immediately upon entering a room.

Close your eyes and try to echolocate your hand while holding it in front of your face. Make hissing noises or repeat a favorite syllable, while moving your hand closer to your mouth and farther away, right and left, up and down. Even better, recruit a friend to move her hand around in front of you. If you're like most of the people we've tried this with, you'll find that you can tell something about where the hand is from the differences in the sound. The hand itself is silent, of course; you're echolocating. We're inclined to think that there's something it's like, phenomenologically speaking, to do this – not just something it's like to move your hand, to make noises, and to hear your own voice, but something it's like to get a sense of where the hand is from hearing the changes in reflected sound as it moves. You have an auditory experience of the hand being very

near or farther away, as moving (maybe) to the right or to the left. You hear, we think, the proximity of your silent hand.

Try another test. Find an empty stretch of floor near a wall. Close your eyes and slowly walk toward the wall, repeating the word “hello”. We venture that you’ll have little trouble stopping a few inches from the wall and that you’ll notice substantial changes in the reflected sound of your voice. When you’re a few inches away, it will *sound* to you like you’re a few inches away. (If you’re concerned that your judgment here will be inappropriately affected by your visual knowledge of where the wall is, have a friend move you an unknown distance from a wall. Analogous changes can also be made to the other tests below.) Next, step from a small room into a larger one, noticing the sound as you walk, then enter a small room again. Just as an orange presents a different visual phenomenology than a grape, so also, we think, does a long hallway present a different echoic phenomenology than a closet. Walking through a doorway, can you hear the frame approach your ear? Sitting at a desk sounds different from sitting in a wide open room, especially if you start talking. Slowly move this book or some other book toward one ear, bringing it within a few inches. You hear the approach of this silent object, don’t you? If unbeknownst to you and outside your visual field, a silent object were approaching your ear in that way, you’d probably react unless you were very absorbed in or distracted by other things – more reason to think that echoic information constantly contributes to your general sense of your environment. Right now, I’d say, it sounds to me like nothing is very near my ears; it sounds like I’m in a clear space.

Try one last test, more pertinent to shape and texture. Closing your eyes and repeating a syllable, slowly move this book toward your face, noticing the changing

sound of your reflected voice. Now do the same with something large and hollow like a mixing bowl or a cardboard box. Now try it again with something soft like a wadded shirt. Don't they sound pretty different? Doesn't the box or bowl sound hollow? Holding the book about a foot from your face, play around with its orientation. Do the same with the box or bowl. Now consider again the opening questions of this chapter.

iv.

Nagel, the philosopher quoted in the epigraph, says that the bat's sonar "is not similar in its operation to any sense that we possess". If the initial reactions of participants in Mike's experiments and reactions from our colleagues in philosophy and psychology are any guide, a significant proportion of the adult population will deny that they can detect, or have any conscious auditory experience of, the size, distance, shape, or texture of silent objects by attending to patterns of reflected sound. I have entertained tables full of philosophers at conference dinners by having them echolocate each other's hands, an activity that reliably produces giggles of surprise even among some of the more sober eminences.

You might think that the blind, whose abilities at echolocation are generally thought to be superior to those of normally sighted people, and who often actively use echolocation to dodge objects in novel environments, would be immune to such ignorance. Not so. For example, one of the two blind participants in Supa and colleagues' 1944 study believed that his ability to avoid collisions was supported by cutaneous sensations in his forehead and that sound was irrelevant and distracted him (p.

144 and 146). Although asked to attend carefully to what allowed him to avoid colliding with silent obstacles, it was only after a long series of experiments, with and without auditory information, and several resultant collisions, that he was finally convinced. Similarly Philip Worchel and Karl Dallenbach (1947) report a nearly blind participant convinced that he detected the presence of objects by feeling pressure on his face. Like Supa's subject, he was disabused of this idea only after long experimentation. (This participant, it turned out, used his impoverished visual sense of light and dark more than tactile *or* echoic information.) Such opinions used to be so common among the blind – until Supa, Dallenbach, and their collaborators demonstrated otherwise – that the blind's ability to avoid objects in novel and changing environments was widely regarded as a tactile or tactile-like “facial vision”, perhaps underwritten by feeling air currents or the like (see Diderot 1749/1916; James 1890/1981; Hayes 1935; Supa et al. 1944; the negligible relevance of air currents is shown by participants' excellent performance when ears are uncovered and cloth is draped over the rest of the face and their poor performance when ears are stopped and the face is left clear). Presumably, if blind people experience auditory echoic phenomenology, and if they are – as people in general are widely assumed to be – accurate judges of their phenomenology, it should occur to them that they detect silent objects at least in part through audition. They should not make such large mistakes about the informational underpinnings of their object sense.

Modus tollens on that last conditional, of course, yields only a disjunction: *Either* blind people don't experience auditory echoic phenomenology (at least not enough to notice its pertinence when prompted to reflect on their remarkable ability to avoid unseen obstacles) *or* they are not accurate judges of their phenomenology. And you might think

the first disjunct considerably more probable than the second. In the journal article on which this chapter is based (Schwitzgebel and Gordon 2000), Mike and I argued that echoic information was unlikely to be experienced phenomenologically as pressure on the face; but now we're not so sure. On the one hand, the cross-modal sensory transformation required seems peculiar: Other than in rare cases of synaesthesia, how often do we experience auditory input as tactile?⁶ In general, background expectations don't seem enough to induce such a change: When you think your cell phone is set to "vibrate" and instead it plays your ringtone, you don't normally (I assume!) experience that ringtone as tactile vibration in your pocket. On the other hand, Mike has come to think he does sometimes experience a feeling of pressure on his face when echoic information is highly salient; and H. Ono and colleagues (1986) found that even when people are aware of the actual informational basis of their abilities, a substantial proportion (both blind and sighted) continues to report experiencing facial pressure. Even if so, we're inclined to think that facial pressure ordinarily does not *replace* auditory phenomenology as much as *accompany* it. If so, those who deny auditory echoic phenomenology are still mistaken about their experience.⁷

This much is straightforward: People often deny the existence of a capacity they demonstrably have, the capacity to detect the position and properties of silent objects via echoic information. The key remaining question – and the question most central to the theme of this book – is whether people err about their sensory experience too. That is, err about not just about their *abilities* but also about their *phenomenology*. Mike and I are inclined to think they do. The phenomenology of echolocation is, or at least often is, auditory; yet people tend to deny that they have such an auditory experience.

Mike and I are wrong in this central contention of ours if one of two things is true: if echolocation does not typically have an auditory phenomenology, or if it does and people generally know that fact. The following thought might seem to support the second of these possibilities. Although people deny that they have auditory experience of silent objects when the question is posed abstractly, that's merely a *theoretical* mistake, not a mistake about their conscious experience. When the matter is put less abstractly – when asked whether they can hear the difference between being in a shower and being in a concert hall or whether they can hear the difference between speaking into a bowl and speaking into a blanket, people will grant that they can hear such differences. They just tend not to think of that as “echolocation” – perhaps because they think of echolocation as an exotic talent of bats and dolphins. It's like the case of timbre, perhaps. Without any background on the topic, someone might deny that he could hear the subtle differences in the overtone series that constitute changes in timbre. While he knows that flutes and trumpets sound different, for example, he could easily be unaware that the physical basis of much of that difference is in the overtone series. He knows his experience of timbre perfectly well, just not under that label or guise. There's no *introspective* error in such a case.

Mike and I acknowledge that there's some merit in this objection. People are not perhaps *as* badly mistaken about their echolocatory phenomenology as it might seem from their tendency toward flat denials when asked abstractly. After all, they know at least that they can hear echoed shouts across canyons and the echoic difference between showers and concert halls. And yet they still are, we think, pretty badly mistaken about their experience. Most people feel surprised – not just about their skills, but also about

their auditory experience – when they discover that they can hear whether a hand or a wall is near their mouth or far away. The parallel does not hold for the person we’ve imagined as doubtful about timbre, who is *not* surprised that he can hear the difference – what we but not he would call the difference in timbre – between a flute and a trumpet. The doubter of timbre experience makes only a theoretical mistake or mistake of labeling – a mistake that is repaired not by introspection but by learning acoustics – while the doubter of echolocation fails to appreciate an introspectively discoverable aspect of her experience.

Is it simply, then, that people are perfectly good at introspecting their echolocatory phenomenology when they consider specific, ongoing instances, such as the hand or bowl in front of the face, and wrong only in their generalizations? No, we think people will tend to be wrong in specific, ongoing instances too. The cases of the hand and the wall and the bowl, described in section iii, are cases meant to be especially striking. In less striking cases, like the case with which Mike and I began this chapter – we began, you’ll recall, asking if you can hear the book in your hands, the desk, the wall – people will tend to be wrong about their current, ongoing auditory phenomenology. You *do* hear such things, we think, at least crudely; they’re part of the rough echoic experience of the objects and space around you, an experience which, if you’re typical, we doubt you appreciated at the time. The echolocatory experience of distance, and other properties, is hard to discern in most cases, despite its pervasiveness. Though it sounds to you like there’s a wall a few feet to your left, and though this sensory knowledge helps guide your behavior, you may have trouble discerning this aspect of your experience even

on careful introspection. You may not – probably did not at the beginning of this chapter, if you're typical – know that you were having such experience, under any label or guise.

This leads us to the other avenue against our thesis: Maybe people don't, in fact, generally have auditory experience of silent objects and their properties. The issue isn't whether people in fact echolocate in their ordinary, daily lives. We addressed that matter as well as we could in section iii, with the anecdotes and thought experiments about shower walls and the Wenger room. The issue, rather, is whether, despite a certain level of sensory attunement to echoic properties, there might actually be no echoic phenomenology. Maybe people only hear sound *sources*, that is, things that emit sound, as modified by environmental conditions, from which they can make *inferences* about silent objects, which remain unheard. Or maybe people have some non-inferential, direct sensory knowledge, but without any experienced sense modality at all – as perhaps is the case when one has the sense that one is being stared at, without knowing how exactly one knows that fact (hearing? peripheral vision? ESP?).

Suppose the philosophy department chair is standing behind me talking and someone tosses a blanket over his head. I'll notice a sudden difference in his voice. But it doesn't seem right to say that I *hear* the blanket. Rather, I hear only his voice, and that it has suddenly changed, from which I infer that something is muffling it. Or if your footsteps are suddenly quieter, you may hear a difference in them that suggests that you've moved from tile to carpet; but perhaps it's not quite right to say that you hear the tile or the carpet. As you approach the wall in an echolocation experiment, you have auditory experience of your own voice saying "hello", and you experience a certain change in the sound of your voice from which you can infer or otherwise learn of the

presence of a wall, but (the objector might say) there is no auditory experience of the wall itself. Such cases are analogous, perhaps, to hearing a change in your dog's barking that indicates that the mail carrier has arrived without hearing the mail carrier herself.

Mike and I find it implausible to draw the sharp line, implicit in the view just sketched, between heard objects that produce sound and unheard objects that reflect or otherwise modify it. People are embedded in massively complex acoustic environments with rich echoic and reverberatory properties, and our sensory systems wisely exploit that fact without, we think, troubling to limit our sensory experience to only sound producers. We find it odd to suppose that we detect this large class of acoustic events yet omit them from our phenomenology. Certainly in vision we experience not just light sources but even more importantly objects that reflect light.

If I clap loudly, the steel bookcase to my left picks up some of that disturbance and rings for about half a second, as I can distinctly hear if I put my ear up next to it. It seems that in that case, at least, I hear the bookcase *as though* it were a sound source – so is it a source, then, in the relevant phenomenological sense? My trashcan rings and reverberates less; my piano more. The reverberations and ringings are more distinctively associated with particular objects if I'm very near them; more generally diffused through the room as I back away. Where's the bright line? And how do we draw the boundaries around a "source", anyway? If I cup my hands around my mouth to shout better, are my hands part of the sound source or are they modifiers of it? If a car crashes through a building, which parts of the car and building produce that horrible noise and which merely modify it? One could, of course, draw a line in the sand or say it's merely vague. But it seems to us that a full appreciation of the complexity of sound environments makes

it unnatural to center one's theory on a robust and principled distinction between sound sources, which people can auditorily experience, and non-sources, which people cannot. Partly for this reason, in arguing for a phenomenology of echolocation, Mike and I are not arguing that the experience of hearing reflective and reverberant objects is different in kind from hearing sound sources; quite the opposite. They are of a piece, all part of our general auditory experience of our surroundings.

I'll grant that I can't hear the blanket that is now covering the department chair's head or even (which is perhaps a weaker claim) *that* he has a blanket over his head. My echolocatory skills are not so finely tuned. I hear only that something soft has interposed. But perhaps someone with longer experience of such things – or a bat – could hear that he has a blanket over his head, as opposed to, say, a rubber mask. Just as, standing in a parking lot, I may auditorily experience an approaching sound source like a car (or would it only be the engine and wheels that are the source?), so also as I step toward a wall – eyes open or eyes closed – I have an auditory experience of the wall's looming nearer. You too: You visually experience depth (or so I'll assume, though see Chapter 2); you auditorily experience the distance of sound sources (for example, a voice may sound far away); and likewise you auditorily experience the distance of silent sound reflecting objects, like the book near your ear or the bowl you're speaking into. Auditory experience of the presence or absence of, and the properties of, reflective and reverberant objects is pervasive. We think that people, hampered by a simplistic folk acoustics and insufficient appreciation of their auditory skills, are inclined to miss this fact when they casually introspect.

Another objection: Maybe even if *Mike and I* experience the world as echoically rich – at least when we explicitly reflect on it – that’s only because our experience is not what it once was. We’re no longer naive about echolocation and thus, maybe, we no longer hear in quite the same way. Theoretical knowledge about echolocation may create echolocatory experiences where none were before; it may create experiences that most people don’t have and thus correctly deny. Mike and I may be erroneously assuming that others are like us, or like us in our most echolocatory-reflective moments. If so, then introspective accuracy can be preserved. Everyone could be right about her own experience.

One reason we’re reluctant to accept this view is that it seems to us that upon noticing echoic properties in our experience – for example, in the hand-in-front-of-the-face demonstration, which was my own first experience with echolocation considered explicitly as such – we have the sense not of creating some new phenomenology but rather of recognizing something familiar, an aspect of experience that had always been there and was only then being made salient and properly considered for the first time. Another reason we’re reluctant to accept the view that echolocation deniers are correct about their own experience is that they sometimes change their minds. They admit, sometimes, in conversations with us, to having been mistaken, not just about their capacities but also about their experience. Recall Supa’s blind participant who regarded sound as “distracting”. The experimenters don’t provide a detailed phenomenological report (despite the fact that Supa’s advisor and co-author Dallenbach had been a student of the great introspective psychologist Titchener; see Chapter 5), but we’d guess that this man, like many of the sighted people we’ve informally interviewed, would have re-

evaluated his auditory phenomenology by the end of the experiment. On the everyone-is-right view, people's confessions of error would themselves have to be mistaken.

People's experience will have had to change radically in character, from nonechoic to echoic, in the course perhaps of a few minutes, while – if they come to think they were previously mistaken about their experience – they failed to notice having undergone this change. This seems an odd pairing of moment-by-moment accuracy with serious ignorance of change over time. One might also wonder whether such a seemingly fundamental shift in experience is likely to be driven by having acquired a little high-level knowledge. Mike and I don't mean to deny that knowledge about echolocation can substantially alter one's auditory experience. Maybe our own auditory experience, now, does differ importantly from that of someone ignorant of human echolocation. All we mean to deny is that the change is so profound as to introduce pervasive echolocatory phenomenology where none was before. Unless there is such a profound change, the echolocation avowers and the echolocation deniers cannot both be right about their experience.

v.

But now that all this has been said and argued and I read back on it, I find myself impelled to make a confession: I'm not actually as confident about my echoic phenomenology as all that. (I can't speak here for Mike, who remains confident.) I recognize that this undercuts the argument of this chapter. So be it. This book is as much confessional epistemology as it is advocacy. The idea of a pervasive echoic

phenomenology does tempt me, both introspectively and theoretically. But as I sit here in my office, I find myself not completely convinced that I really do have auditory experience of the wall to my left and the desk before me, as I tap on my keyboard and the hum of freeway traffic penetrates my window. I close my eyes and speak aloud. I feel pretty confident that I can hear that my voice isn't being reflected back by an object within three inches; but do I really hear, even crudely, the various silent things around me? When I went to Rosenblum's lab and tried out the Rosenblum-Robart shape detection experiment for myself, I did very well. In a short run of ten trials, I got nine right. Evidently, I could hear whether there was a triangle, circle, or square before me. So now I close my eyes, and saying "da da da da" I lean in close to my computer monitor. Can I hear that it's square? Even if I could rightly guess that it's square, just from the hearing, is there some distinctive auditory phenomenology of squareness, analogous to the visual phenomenology of squareness that I experience with my eyes open?

Well, now you have as many tools to judge in your case as I have to judge in mine.

¹ Nagel, by the way, is a philosopher of mind with no particular background in perceptual psychology. He offers this claim simply as an obvious, commonsensical point. Nagel acknowledges in a footnote on a later page that blind people can use something like bat sonar when tapping with a cane. He appears to have trouble imagining even what that is like, writing in the subjunctive, “Perhaps if one knew what that [i.e., the sonar of the blind] was like, one could by extension imagine roughly what it was like to possess the much more refined sonar of a bat” (1974, p. 442).

² It is I who stress the specifically *auditory* character of echolocation. Mike writes:

I am more doubtful than Eric seems to be about whether echolocation, or really any form of perception, can be phenomenologically or even biologically separated from the other modalities (barring organisms – paramecia? – with only a single sensory system). Humans, bats, and dolphins all echolocate, *and* will use visual, tactile, vestibular, olfactory, and any number of other signals as they become available. Consider that the bat does not detect its echolocatory capacities; rather echolocation is a means by which it can detect an insect flying past. To suggest that the phenomenology of the bat is different in kind from any other species is to suggest that the principal determinant of phenomenology is the medium by which sensory information is detected – and not the content of that information.

Furthermore, while in this chapter we have outlined some of the literature to suggest that echolocation is a valid sensory channel for both blind and sighted humans, in a larger sense one could argue that all of our perception is part of a

unified multisensory experience (including various forms of hearing) that informs us about the world and the objects therein. Did perception evolve to inform us about how we detect the world, or to inform us about the things in the world that will support successful actions? While most perceptual psychologists still embrace the classic Aristotelian division of the sensory systems, this perspective has caused me to favor a more continuous amodal/modality-neutral approach that distinguishes sensory experiences by their organization in the environment (i.e., the object and events occurring in the world), and not by the sensory channels by which they are detected.

The case for error about experience becomes somewhat more complicated on the modality-neutral approach, since people are more often right, in my view, about the fact that certain environmental objects are sensorily present than they are about the specific sensory character of their experience of those objects. However, even on the modality-neutral approach, people are still mistaken when they deny having *any* sensory experience of the wall, the book, etc., with their eyes closed.

³ Traditionally, echolocation was defined only “monostatically”, as the detection of the reflections of self-generated sounds, as in Griffin 1958 and Simmons 1989. However, some researchers have also emphasized the importance of sounds generated from other sources, as in Lee 1990; Xitco and Roitblat 1996; and Rosenblum et al. 2000. Since the cognitive and functional demands seem to be similar in the two cases, we opt for the broader definition.

⁴ Prior (1999) describes training a dolphin to use eye covers (a dolphin blindfold) for a marine show intended to exhibit the dolphin’s excellent echolocation abilities.

While the perceptual task was easy and the eye suction cups were designed not to be physically irritating, the dolphin was resistant to them. Through a series of ingenious training methods, Prior managed to convince the dolphin to accept the eye covers, but the dolphin's resistance is interesting. How much do dolphins know about their sensory abilities? Experiments by Gopnik and Graf (1988) and Oneill et al. (1992) suggest that three-year-old children often don't know very well the sensory bases of their knowledge or what information comes through what sensory modality. Perhaps dolphins have the same difficulty? Or maybe this dolphin just didn't want to sacrifice his visual input, despite the safe environment and easy task?

⁵ Rosenblum and Robart (2007) found that sound intensity (measured in decibels) was virtually identical for the triangle, circle, and square when measured at listeners' average head position, and listeners were still able to detect the shape differences at rates above chance. Gordon and Rosenblum (2004) tested the effects of intensity by randomly varying, along with the aperture size, the sound intensity of the crowd noise from the speakers. The sound at the listener's average head position was always substantially quieter in lowest sound intensity condition, regardless of aperture size, than it was in the highest sound intensity condition, so that if participants were judging strictly by intensity they would judge all the apertures more passable in the high sound intensity condition than in the low sound intensity condition. Gordon and Rosenblum found, on the contrary, that although sound intensity did have an effect on judgments of passability, with apertures more likely to be judged passable at lower sound intensities, the larger apertures were much more likely to be judged passable than the smaller apertures, regardless of sound intensity. For example, the largest aperture was judged passable 69%

of the time at the lowest sound intensity while the smallest aperture was judged passable only 47% of the time at the highest sound intensity.

⁶ Sometimes we experience vibration both tactilely and auditorily, like when we feel the buzz of a loud, thumping bass line, but that is not so much auditory input being experienced tactilely as a single event that experienced simultaneously via two different senses. It's like both seeing and feeling a square. On synaesthesia, see Simner et al. 2006; Hochel and Milán 2008.

⁷ Still another view is that of Kells (2001), who extensively interviewed eight blind people about the phenomenological character of their ability to detect silent objects. She characterizes the experience neither as auditory nor as involving facial pressure but rather as a sensation of its own unique kind – a hard-to-convey sense of the presence of something in space or of openness or closedness. Not all blind echolocators seem to talk this way – not, for example, Kish (2009), nor the participants in Worchel and Dallenbach 1947 who often denied having the ability to detect objects at all in the echolocatory tasks, despite pretty good performance).