

How Well Do We Know Our Own Conscious Experience? The Case of Human Echolocation

Eric Schwitzgebel
University of California, Riverside

Michael S. Gordon
University of California, Riverside

In “What Is It Like to Be a Bat?” the philosopher Thomas Nagel says that “bat sonar [or echolocation], 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.”¹ We do not know what “it is like” to echolocate, he claims, and so we don’t know what it is like to be a bat. We agree that most people know very little about what it is like to echolocate. However, such ignorance entails neither the inability to echolocate nor the absence of a conscious experience of echolocation.

We will argue, first, that sighted human beings do echolocate as part of their normal intercourse with the world. We will then argue that although there is something it is like for us to echolocate, most people’s knowledge of that experience is very poor indeed. We think, contrary to what most philosophers seem to suppose, that normal people in normal circumstances can be grossly and systematically mistaken about their own current conscious experience.

EMPIRICAL SUPPORT FOR THE CLAIM THAT NORMAL HUMAN BEINGS ECHOLOCATE

Echolocation, as we will be using the term, is the ability to detect the reflective and reverberant characteristics of an object or an environment using sound generated in the area.² It is well known that several mammals (e.g., bats, dolphins, whales) can echolocate with great accuracy.³ It has also been found that blind people can use this ability as a primary source for navigation while they walk (or even bicycle) in unfamiliar environments, and sighted individuals wearing a blindfold can echolocate effectively with brief training.⁴

For example, Michael Supa and his colleagues instructed blind and sighted subjects (wearing a blindfold) to walk toward a 4' x 4'10" masonite board mounted at various distances from the subjects' starting position.⁵ During their approach, subjects were asked to signal the moment of their initial detection of the board and also when they were as close to the board as possible without touching it. Blind subjects were able to detect the board several feet before contact and move within a few inches of its surface—nearly always avoiding contact. The sighted subjects achieved a similar level of accuracy after about thirty trials.

To determine the information supporting this ability, subjects were asked to walk toward the board wearing socks, but no shoes, to reduce the amount of generated noise. Following this manipulation, the subjects' hearing was effectively eliminated using earstops and, later, projection of noise directly into their ears. Performance in these conditions degenerated proportionately with the amount of induced deafness to the point where all subjects collided with the board every trial.

Lawrence D. Rosenblum and Carol H. Ammons and their respective colleagues found similar results.⁶ Sighted but blindfolded subjects, after brief training, were able to stop the moment before walking into a large, sound-reflecting surface. Rosenblum and his colleagues also found that sighted subjects were able to use echolocation to distinguish between walls distanced 36 to 144 inches in front of them. In this study, blindfolded subjects echolocated a wall and then, with the wall removed, estimated its distance by walking to where they thought it had been. Sighted individuals were able to detect 36-inch differences between wall positions.

Steven Hausfeld and his colleagues asked blindfolded, sighted subjects to echolocate a target placed twenty-five centimeters in front of the face.⁷ Targets were varied in shape (circle, triangle, or square, all of equal surface area, or no target) or texture (fabric, plexiglass, carpet, or wood). After very brief training, the subjects were able to determine which of the shapes had been presented and to distinguish between some of the textures. Charles E. Rice also suggests that both blind and sighted people can discern the shape and location of objects by echolocation.⁸

This echolocation ability may be due, in part, to several different sources of information. If a listener produces sound, she may be able to use the time delay and loudness differences between production and reception to determine a reflective surface's distance. Reflected sound from any source potentially can provide a listener with distance and object property information using timbre and pitch variance. A moving listener may also use timbre changes (caused by the interference of sound reflecting back on itself), motion parallax information (closer objects change their position relative to an observer much faster than distant ones), tau (objects expand much more rapidly as they get closer), and Doppler shift.⁹ Such auditory information is available to both blind and sighted individuals and may provide the informational support for the abilities demonstrated throughout this echolocation research.

It does not follow from the fact that sighted individuals can be trained to echolocate in experimental settings that they actually do so in their daily lives. Nonetheless, we do think it is plausible to endorse this stronger conclusion. Consider, for example, your own experience in walking down a long, tiled hallway. With each step, a small burst of noise radiates from your shoe's contact with the floor into the area around you. In hearing this, 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 acoustic experience would be noticeably different. For example, there is an obvious echoic difference that makes you sound like Pavarotti in the shower, and you everywhere else. Hallways and showers sound different. They do so not because hallways and showers *produce* sound but because they *reflect* sound differently. In reacting to the acoustic differences between the two, one is using echoic information.

While our primary mode of perceiving in our daily, sighted lives, is not echolocatory, we do constantly seem to use echoic information as a supplement to vision. If what we think to be a concert hall does not sound like a concert hall or what we think to be a hallway does not sound like a hallway, we will detect that difference. Imagine stepping through a doorway into a familiar tile hallway, visually focusing to the left, and being surprised by the sound of one's footstep. Turning right, one discovers a large piece of furniture where there was none before. In such a case, echoic information has been used to guide visual attention. While moving through the world, we constantly use echolocatory information to assess our physical surroundings and determine the appropriate behavioral responses.

Wenger Corporation has developed what they call a "virtual room." This room is able to synthesize the acoustics of a variety of spaces, from an office to a symphony hall (something that practicing musicians have found quite useful). If the acoustics of the virtual room are set to emulate an area

much larger than the actual size of the room, listeners can quickly determine that something is amiss. Typically, individuals entering the room will immediately glance upward to confirm or disconfirm their acoustic apprehension of the room. If echoic information were not being regularly used to supplement other sources, one would not expect such a reaction. Clearly, one does not normally glance at the ceiling immediately upon entering a room.

Human beings, not just bats and dolphins, have a sense of the acoustic space they occupy. Sounds generated both by objects in the environment and by ourselves, combined with echoic reflections off the surrounding surfaces, facilitate a supramodal perceptual experience of the world around us. Wenger's virtual room manipulates this sense of space, causing what appears to be a powerful, echolocation-guided reaction. More research is needed, however, to determine precisely what sorts of echoic information sighted people respond to and how capable we are in using it.

THE POOR QUALITY OF OUR KNOWLEDGE OF THE PHENOMENOLOGY OF ECHOLOCATION

It seems to us that there is something it is like to echolocate, i.e., that there is a conscious phenomenology of echolocation, that just as an orange presents a different visual phenomenology than a grape, so also a hallway presents a different echoic phenomenology than a shower. If we are wrong about this matter, that only makes more plausible our ultimate point: that people can be vastly wrong about their own conscious experience.

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, and move your hand closer to your mouth and farther away, right and left, up and down (or even better, get a friend to move her hand in this way). In doing this, you can tell something about where your hand is from the differences in sound; you are echolocating. We hope you'll agree that there is something it is like to do that—not just something it is like to move your hand, to make noises, and to hear your own voice, but something it is like to get a sense of where your hand is from the changes in the reflected sound as it moves. You have an auditory experience of your hand as being closer or farther away, to the right or to the left. Although not as vivid as visual experience, echolocation is an important, pervasive, and distinctive feature of our sensory phenomenology.

It is interesting, then, to see how grossly mistaken people can be about it. Nagel, as quoted at the beginning of this paper, goes so far as to claim not only that we do not echolocate as the bat does, but also that the bat's sonar "is not similar in its operation to any sense that we possess."¹⁰ Maybe all Nagel means is that the bat's echolocation is so vastly better than ours as to warrant description as different in kind. Yet, Nagel puts his point more

strongly than would seem to be warranted if that were his view. And Nagel is not alone in his apparent failure to recognize the human capacity for echolocation. If the initial reactions of subjects in the second author's experiments¹¹ and reactions from colleagues in cognitive science and philosophy can be used as a guide, a significant proportion of the adult population will deny that we can detect object properties (e.g., size, shape, texture, distance) by attending to patterns of reflected sound.

One might think that the blind, whose abilities at echolocation are generally considered superior to those of sighted people, and who often use echolocation quite actively to avoid objects in novel environments, would be immune to such errors, but that is not the case. For example, one of the two blind subjects in Supa and his colleagues' experiments believed that his ability to avoid collisions with objects was supported by cutaneous sensations in his forehead and that sound was irrelevant and "distracted" him.¹² Although asked to attend carefully to what allowed him to avoid colliding with obstacles, it was only after a long series of experiments, with and without auditory information, and several resultant collisions, that he was finally convinced that his judgments were based on auditory perception. Similarly, Philip Worchel and Karl M. Dallenbach report a nearly blind subject convinced that he detected the presence of objects by pressures on his face. Like the subject in Supa et al., he was disabused of this idea only after long and varied experimentation.¹³ In fact, so common was it—until Dallenbach and his collaborators demonstrated otherwise—for blind people to think that they detected objects by feeling pressures on their faces, rather than by echolocation, that their ability was originally called "facial vision."¹⁴ It is striking that such an error could occur despite the likelihood that a good number of blind people directed extended attention to what in their perceptual experience allowed them the remarkable ability to avoid obstacles they could not see.

Let's stop to consider exactly what sort of error is being made by people who deny the experience of echolocation. Descartes is reputed (with what justice we are uncertain) to have thought that although we may be radically mistaken in our judgments about things external to our minds, we cannot be mistaken in reporting our own mental states (or at least our own current conscious experiences if we reflect attentively upon them). It may be that a person who believes she is staring at a red tomato is not staring at a red tomato at all—she may, instead, be having a drug-induced hallucination of a tomato, or she may have had her brain removed the previous night and teleported to a vat on Alpha Centauri where evil neuroscientists are stimulating it in such a way as to mimic exactly the experience of seeing a red tomato. Although these scenarios may be unlikely, it is often granted that they are possible. She may be mistaken in thinking that she is seeing a red tomato. However, there is one thing about which it might appear she

cannot be mistaken: that she is having a subjective or phenomenal experience of “redness.”

Many philosophers since have found plausible a “Cartesian” view of this sort: We cannot be mistaken about, or be brought sincerely to doubt, the subjective quality of our own current, conscious experience.¹⁵ Although recently this view has fallen into disfavor, the philosophers who repudiate it typically seem to suppose that our knowledge of our current conscious experiences is nonetheless quite accurate.¹⁶ They often rely on far-fetched, science-fiction examples involving scientists with brain-o-scopes, or examples of mistakes made only for a moment, or mistakes made at the margins of our discriminative abilities. The dialogue, in other words, has been mostly between those who regard introspection of conscious experience as infallible and those who regard it as fallible in principle but nonetheless reliable (exceptions to this trend include the Churchlands and Hilary Kornblith, but even their examples of errors tend to be marginal or pathological).¹⁷ *Prima facie*, denial of the experience of echolocation would seem to be a gross and pervasive error of exactly the sort disallowed by both of these opposing viewpoints. Could people really be making that big a mistake about their own phenomenology?

We think it plausible to say that, at least in some cases, people are making such a big mistake. However, let’s examine a few of the more appealing alternative hypotheses that might come to mind. We will argue that they all have significant weaknesses.¹⁸

First, let’s consider the possibility that the phenomenon could be explained wholly by a failure of people to attend to their own experience. Perhaps people deny echolocating simply because they have never bothered to focus on that aspect of audition. If so, then a certain sort of Cartesian view, one in which the perceiver is accurate when he actively attends to his phenomenology, could be supported. This proposal gains some plausibility from the fact that certain tests (e.g., the hand-in-front-of-the-face demonstration) that draw our attention to the echolocatory aspects of auditory experience—at least when presented as tests of echolocation—do cause most people to recognize an echolocation experience.

We grant that our ignorance of echolocation, especially in the case of normal, sighted subjects, has a lot to do with not carefully attending to patterns of change in our auditory phenomenology. Once one does start regularly attending to the acoustic changes that occur within a normal echoic environment, the experience of echolocation starts to seem both obvious and pervasive. However, for several reasons we do not think failure to attend can wholly explain our ignorance of echolocation. First, it is rather odd to suppose that most people never attend to a major and persistent feature of their auditory experience (especially if one is willing to grant—as we are not—another view associated with Descartes, that it is by knowledge of

our experience that we gain knowledge of the world outside of us). Second, blind subjects who claimed to detect the presence of obstacles by facial vision rather than by echolocation *were* attending to their auditory experience, presumably alert for changes in it that might signal the approach of an obstacle, and nonetheless continued to make the mistake. Finally, these same subjects not only claim not to echolocate, but also claim to feel pressure on their faces. (Several of the sighted subjects in the Rosenblum, Gordon, and Jarquin study claimed a similar thing when questioned about their experiences.)¹⁹ This is not a failure to notice something to which they are not attending, but rather a confabulation of something new, and thus requires a different explanation.

Another possibility is that our mistakes about echolocation are just an error of words: It may be that people can, when attending, recognize in themselves the experience of echolocation, and they deny that the term *echolocation* applies either because they do not understand the term or because they do not know enough about the physical causes of their experience to consider it appropriate. Certainly some errors that might seem on the surface to be mistakes about one's phenomenology are really verbal mistakes instead. If I do not know to what color the word *maroon* refers, I might mistakenly deny that I am having an experience as of seeing a maroon object. In so doing, I am making a mistake in *reporting* my phenomenology, but I am not mistaken about what the experience itself is like. My confusion is only a confusion of words.

However, the mistake of the naïve subject who denies that she is echolocating cannot be a simple mistake of this sort. She not only will refuse to accept the term *echolocation* to describe her experience, but also will deny that she can detect the properties of silent objects by attending to the sounds they reflect when that claim is couched in language she clearly understands. If her mistake is a verbal mistake of reportage, a more complex story needs to be told.

For example, a person may be able to recognize the difference between the experience of hearing a middle-C played with one timbre (characteristic of a flute, say) and another middle-C similar in all respects except for the timbre (characteristic of a trumpet). This person may not know the word *timbre* and may find implausible the idea that he could be sensitive to those sorts of differences in the overtone series. In denying that he is sensitive to timbre, he is making a mistake about what stimuli in the environment move him to have different experiences, but he is not making a mistake about his experiences themselves. He knows darn well the difference in auditory experience between the two tones.

However, the echolocation case is not like that. The echolocating subject does *not* recognize that there are different auditory experiences resulting from the distinct reflective patterns of objects divergent in shape or

distance. When asked whether the experience of hearing a flute is any different from the experience of hearing a trumpet, our skeptic about timbre will say yes; but he will deny the analogous difference in the echolocation case. As discussed above, this is not simply a matter of failing to attend, since subjects make these mistakes even when attending to their auditory experiences and attempting to assess whether there are any changes in them that might serve as the basis for distance judgments. The person who fails to notice significant features of her sensory experience (and perhaps even confabulates other features such as facial pressure) is not merely making a verbal mistake or a mistake about the physical causes of her experience. She is making an error about the experience itself.

Another alternative is to deny that there is any conscious experience of echolocation. The phenomenology of audition, on this view, might be a raw buzz of sense data, unaltered by any sort of higher-level processing about (say) the positions of objects, or it might incorporate higher-level information, but not echolocatory information—at least not in such a way that it is right to say that there is a phenomenology of echolocation. Our echolocatory capacities, on this view, are the consequence of unconscious inferences our perceptual system makes on the basis of auditory input but does not incorporate into the sensory experience itself. Thus, subjects who report no experience of echolocation are reporting their phenomenology accurately.

Such a view, however, may not help the reliabilist about introspection. Although it renders accurate those who deny any experience of echolocation, it correspondingly renders inaccurate those who claim that there is such an experience. Unless the reliabilist can explain away the latter assertion as merely an error of words (discussed above), such an error would seem to be an error of the sort the reliabilist disallows.

In any case, this view draws an unlikely picture of our sensory experience. We do not, most philosophers and psychologists would now say, experience the world as raw, unprocessed sense data on the basis of which we make inferences about how things are. Rather, higher-level processes affect the sensory experiences themselves. The senses present the world *as being a certain way*. The experience of seeing a blue book, for example, is not just the experience of having a blue sense datum in part of one's visual field; it is the experience of seeing some object as being in some location. Illusions and ambiguous figures show this point nicely. We experience the ambiguous figure as a drawing of an old or a young woman, or we experience the illusion of depth in a two-dimensional figure, because those figures are visually experienced not simply as black and white blobs but as parts of a rich world of objects. Auditory experience, too, is not just raw frequency and intensity but the experience of something as having specific properties (e.g., a certain distance and direction from the observer). Once we are familiar with the phenomenon of echolocation, it seems odd to say that we have an

auditory experience of the distance and direction of sound sources only and not of silent objects. If your ear is near a wall or your hand is near your mouth, part of your auditory experience is of something nearby in that space. Failing to notice echolocation is simply failing to notice that part of your auditory experience.

Furthermore, like both of the previous explanations, this explanation does not address the apparent error of those who claim to experience pressure.

A final possibility is that people consistently report their experience accurately, despite their varying reports. People feel exactly what they report they feel. Those who say that they do not experience echolocation do not in fact experience it; those who say they experience pressures on their faces do experience such pressures, even though the physical cause of those phenomena is stimulation of the cochlea. This hypothesis leads to a number of difficulties.

First, although perceivers may subjectively experience a stimulus in a modality other than the one in which the stimulus was initially presented (a kind of synesthesia), such occurrences appear to be unusual.²⁰ Furthermore, it is unclear why some blind people would experience pressure on the face in the looming presence of a wall while others report auditory sensations. At least two blind subjects became convinced as a consequence of experiments cited above that the alleged facial pressures were imaginary.²¹ Unless they judged the pressures nevertheless to be genuinely experienced, infallibility in reporting sensory experience would imply that their experiences must somehow have changed sensory modality during the course of the experiments. Echolocation in sighted people obviously creates analogous difficulties. Why should some sighted people experience echolocation and others fail to, when it seems unlikely that there is any gross neurophysiological difference between the two groups? It is also suspicious that people can pass from claiming they do not experience echolocation to claiming that they do in a matter of minutes. If such claims reflect a genuine change in opinion about the experiences, then the synesthetic effect must dissipate quite rapidly. But it seems odd to suppose that the sensations can so quickly be rerouted from one modality to another. Are we to suppose that subjects spend their whole lives experiencing echolocation not at all, or as pressure, and then after some bit of formal experimentation shift to experiencing it in a dramatically different way? What would cause such a radical change in sensory experience? (And what would the reliabilist make of a case in which the purported change went completely unnoticed, the subject reporting, instead, that she had been mistaken in her previous assessment and had auditory echolocation experiences all along?)

It may be tempting to defend the reliabilist's view here by suggesting that differences in expectations are responsible for the differences in

experience—i.e., that those who expect to feel pressure will, in consequence, experience pressure and that those who expect auditory experiences will have those instead. There is some plausibility to the hypothesis that expectations can substantially affect one's sensory experience. But even if we grant that the expectation, say, of hearing a child cry in the distance can affect one's auditory experience so that it seems as though a cry has occurred, or the expectation of a certain flavor in the wine subtly changes how the wine tastes, the consistent and regular transformation of sensory experience from one sensory modality to another would seem to be a rather different matter. Pagers, for example, can often be set either to vibrate or to ring when a call comes in. When someone expects his pager to vibrate and instead it rings, he does not normally experience that ringing as a vibration. Yet something like that would have to happen for expectations of facial pressure to cause the transformation of auditory input into a haptic experience.

We are not saying that there is no way to think of Nagel's mistake as something other than a mistake about current conscious experience. However, we hope to have shown that the simplest, most obvious, alternative explanations have significant shortcomings. The most straightforward remaining explanation is that people in Nagel's position are seriously mistaken about their own current conscious experience. We see no compelling reason not to accept that explanation.

CONCLUSION

In summary, then, we contend that human beings not only can echolocate if trained, but actually do echolocate in their daily lives. Furthermore, although there is something it is like to echolocate, most people have poor knowledge of what that is. People make basic mistakes about this aspect of their own conscious experience. If so, then the Cartesian view that we cannot be mistaken about our own current experiences is false—and not false just by virtue of marginal cases involving clinical pathology or evil futuristic neuroscientists, but erroneous with respect to a significant aspect of our daily experience.

In fact, like Kornblith, we are tempted to turn the Cartesian view on its head: What we know most directly and certainly is the external world of middle-sized dry goods; what we know only by uncertain and unreliable inference are our own mental states, including our current, conscious phenomena.²² Not only are we pervasively mistaken about the experience of echolocation, but we make vast and frequent mistakes about our sensory experience in other modalities, about our moods and emotions, about our imagery and dreams, about our cognition and desires.²³ Even though we echolocate, Nagel may still be right that we cannot know what it is like to be a bat; we hardly even know what it is like to be ourselves.

NOTES

We would like to thank Denise Durance, Dillon Emerick, Bryan Lee, Shaun Nichols, Pauline Price, Larry Rosenblum, David Rosenthal, Josh Rust, and our colleagues in the U.C. Riverside Department of Philosophy for discussions that helped to shape this paper.

1. Thomas Nagel, "What Is It Like to Be a Bat?" *Philosophical Review* 63 (1974): 438.
2. Traditionally, echolocation has been considered only "monostatically," as the detection of the reflections of self-generated sounds, as in Donald Griffin, *Listening in the Dark* (New Haven, Conn.: Yale University Press, 1958), and James A. Simmons, "A View of the World through a Bat's Ear," *Cognition* 33 (1989): 155–99. However, some investigators have emphasized also the importance of reflections of sounds generated by third parties; for example, David N. Lee, "Getting Around with Light or Sound," in *Perception and Control of Self-Motion*, ed. Rik Warren and Alexander H. Wertheim (Hillsdale, N.J.: Lawrence Erlbaum, 1990): 487–505; Lawrence D. Rosenblum, Michael S. Gordon, and Luis Jarquin, "Echolocating Distance by Moving and Stationary Listeners," *Ecological Psychology* 12 (2000): 181–206; Simmons, "The Role of Auditory Scenes and Object Flow in Echolocation," presented at the Tenth International Conference on Perception and Action, Edinburgh, August 12, 1999; Mark J. Xitco and Herbert L. Roitblat, "Object Recognition through Eavesdropping: Passive Echolocation in Bottlenose Dolphins," *Animal Learning and Behavior* 24 (1996): 355–65.
3. For example, see W. E. Evans, "Echolocation by Marine Delphenids and One Species of Fresh-Water Dolphin," *Journal of the Acoustical Society of America* 54 (1973): 191–99; Griffin, *Listening in the Dark*; W. N. Kellogg, "Echo Ranging in the Porpoise," *Science* 128 (1958): 982–88.
4. For example, see Milton Cotzin and Karl M. Dallenbach, "'Facial Vision': The Role of Pitch and Loudness in the Perception of Obstacles by the Blind," *American Journal of Psychology* 63 (1950): 485–515; B. McCarty and P. Worchel, "Rate of Motion and Object Perception in the Blind," *New Outlook for the Blind* 48 (1954): 316–22; Charles E. Rice, "Human Echo Perception," *Science* 155 (1967): 656–64; Rosenblum, Gordon, and Jarquin, "Echolocating Distance by Moving and Stationary Listeners"; Michael Supa, Milton Cotzin, and Karl M. Dallenbach, "'Facial Vision': The Perception of Obstacles by the Blind," *American Journal of Psychology* 62 (1944): 133–83.
5. Supa, Cotzin, and Dallenbach, "'Facial Vision': The Perception of Obstacles by the Blind."
6. Rosenblum, Gordon, and Jarquin, "Echolocating Distance by Moving and Stationary Listeners"; Carol H. Ammons, Philip Worchel, and Karl M. Dallenbach, "'Facial Vision': The Perception of Obstacles out of Doors by Blindfolded and Blindfolded-Deafened Subjects," *American Journal of Psychology* 66 (1953): 519–54.
7. Steven Hausfeld, Roderick P. Power, Angela Gorta, and Patricia Harris, "Echo Perception of Shape and Texture by Sighted Subjects," *Perceptual and Motor Skills* 55 (1982): 623–32. We have informally replicated the results of their shape discrimination experiment using a square, a triangle, and two rectangles of different dimensions.
8. Rice, "Human Echo Perception."
9. For a fuller review of information sources possibly used in echolocation, see Rosenblum, Gordon, and Jarquin, "Echolocating Distance by Moving and Stationary Listeners"; Thomas A. Stoffregen and John B. Pittenger, "Human Echolocation as a Basic Form of Perception and Action," *Ecological Psychology* 7 (1995): 181–216.
10. Nagel, "What Is It Like to Be a Bat?" 438.
11. Rosenblum, Gordon, and Jarquin, "Echolocating Distance by Moving and Stationary Listeners."

12. Supa, Cotzin, and Dallenbach, "'Facial Vision': The Perception of Obstacles by the Blind," 146, 144.
13. Philip Worchel and Karl M. Dallenbach, "'Facial Vision': Perception of Obstacles by the Deaf-Blind," *American Journal of Psychology* 60 (1947): 502–53.
14. See, for example, Samuel P. Hayes, *Facial Vision or the Sense of Obstacles* (Watertown, Mass.: Perkins, 1935) and Supa, Cotzin, and Dallenbach, "'Facial Vision': The Perception of Obstacles by the Blind." Both these works describe the long debate about the means by which blind people can sense obstacles. See also Denis Diderot, "Letter on the Blind for the Use of Those Who See," in *Diderot's Early Philosophical Works*, ed. M. Jourdain (Chicago: Open Court, 1916).
15. For example, A. C. Ewing, *The Fundamental Questions of Philosophy* (London: Routledge, 1951); C. I. Lewis, *An Analysis of Knowledge and Valuation* (La Salle, Ill.: Open Court, 1946); Sydney Shoemaker, *Self-Knowledge and Self-Identity* (Ithaca, N.Y.: Cornell University Press, 1963).
16. For example, D. M. Armstrong, "Is Introspective Knowledge Incorrigible?" *Philosophical Review* 72 (1963): 417–32; Donald Davidson, "First-Person Authority," *Dialectica* 38 (1984): 101–11; Christopher Hill, *Sensations* (Cambridge: Cambridge University Press, 1991).
17. Paul M. Churchland, *Matter and Consciousness*, rev. ed. (Cambridge, Mass.: MIT Press, 1988); Hilary Kornblith, "What Is It Like to Be Me?" *Australasian Journal of Philosophy* 76 (1998): 48–60. Patricia S. Churchland even briefly mentions echolocation by the blind in this connection in "Consciousness: The Transmutation of a Concept," *Pacific Philosophical Quarterly* 64 (1983): 81–82. Daniel C. Dennett, in *Consciousness Explained* (Boston: Little, Brown, and Co., 1991), may in some ways be a closer ally, since he discusses normal cases in which there appear to be mistaken reports about conscious experience. However, we find his ultimate position on conscious experience somewhat difficult to interpret.
18. The first hypothesis and something like the third hypothesis were suggested to us by Shaun Nichols in email correspondence, although they may not represent his final view.
19. Rosenblum, Gordon, and Jarquin, "Echolocating Distance by Moving and Stationary Listeners."
20. If one considers McGurk effects and similar phenomena to be synesthetic (Harry McGurk and John McDonald, "Hearing Lips and Seeing Voices," *Nature* 264 [1976]: 746–48), then there is a type of synesthesia that is probably rather common. However, such sensory modality interactions are quite a different sort of phenomenon from the replacement of a sensation in one modality with that in another as envisioned by the hypothetical framer of the objection under discussion. However, see J. Kevin O'Regan and Alva Noë, "A Sensorimotor Account of Vision and Visual Consciousness," forthcoming in *Behavioral and Brain Sciences*, for a discussion of synesthesia and echolocation. O'Regan and Noë address neither the variability of subjective reports about the experience of echolocation nor echolocation in sighted people.
21. Supa, Cotzin, and Dallenbach, "'Facial Vision': The Perception of Obstacles by the Blind"; Worchel and Dallenbach, "'Facial Vision': Perception of Obstacles by the Deaf-Blind."
22. Kornblith, "What Is It Like to Be Me?"
23. See Schwitzgebel, "How Well Do We Know Our Own Conscious Experience? The Case of Visual Imagery" and "Why Did We Think We Dreamed in Black and White?" (both under review) for further developments of this theme.