

PERCEIVING THE LOCATIONS OF SOUNDS

European Review of Philosophy, 7, 2007

Casey O'Callaghan

Bates College

Abstract

Frequently, we learn of the locations of things and events in our environment by means of hearing. Hearing, I argue, is a locational mode of perceiving with a robustly spatial phenomenology. I defend three proposals. First, audition furnishes one with information about the locations of things and happenings in one's environment because auditory experience itself has spatial content--auditory experience involves awareness of space. Second, we hear the locations of things and events by or in hearing the locations of their sounds. Third, we auditorily experience sounds themselves as having relatively stable distal locations. I reject skepticism about spatial audition and auditory experience tracing to Strawson's *Individuals*, and suggest that spatial auditory experience grounds a form of perceptual access to objects and events that is critical to negotiating one's environment.

1. Space and audition

When you hear the sound of a car that is driving by on the street outside your window, you learn not only whether the car has a hole in its muffler or has squealing brakes. You also learn something about the location of the car because hearing furnishes information about the locations of its objects. By listening, you learn not only about the character of the things and happenings around you, but also about where they are in the surrounding environment.

The focus of this essay is the way in which hearing furnishes one with spatial information about one's environment. Hearing, I shall argue, like vision and probably unlike olfaction, is a locational mode of perceiving. When you hear the sound of a car passing or hear the sound of a plate breaking on the kitchen floor behind you, you hear not just something about what is going on, but also something about where it occurs. But hearing does not just provide information that enables subjects to work out or to infer the locations of sound sources. And it is not just a matter of some subconscious process that results in a subject's knowing about where things are. Auditory experience, I want to argue, represents space and locations. Conscious hearing presents the world as being a certain way with respect to space and locations, in addition to its qualitative features and particulars. Auditory experience may thus be more or less veridical or illusory with respect to spatial attributes of one's environment. If correct, this claim illuminates why audition is so useful in negotiating a world populated with things and happenings. Since

it is capable of representing space and, in particular, the locations of sound sources, conscious audition provides the means for an object-involving mode of perceptual access distinct from that of vision.

It is one thing to say that hearing and auditory experience are spatial, and it is another to characterize the experience of auditory space. Here, too, I have a proposal. I shall argue that auditory experience represents sounds themselves as located. In particular, I shall argue for the phenomenological claim that sounds seem to normal human hearers in a wide range of circumstances to be located at a distance and in a particular direction. Sounds seem to occupy relatively stable distal locations, and do not except in special situations seem to travel, to come from their sources in a spatial sense, to be near the ears of subjects, or to fill spaces completely.

My claim about the phenomenology of locational hearing--about how auditory experience represents sounds in space--is a claim about the particular way in which auditory awareness is spatial. It serves as a proposal to explain why we form beliefs and how we learn about the locations of ordinary things and events on the strength of hearing. We learn about the car's location through hearing because the experience of hearing conveys information about the car's spatial location. In particular, hearing conveys such information because we auditorily experience the car's sound as taking place at a certain location in egocentric space.

My plan is this. I shall argue that empirical evidence supports the view that auditory experience is spatial (Section 2) and that it represents sounds as located at some distance in a particular direction (Section 3). I consider and respond to several objections that grant the spatiality of audition while resisting my phenomenological claim about the locatedness of sounds (Sections 4 and 5). I defend the claim that hearing furnishes information about the locations of things and events in one's environment by or in presenting sounds as located. In the remainder of the paper (Sections 6 through 9), I take up an influential group of challenges to the claim that auditory experience has spatial characteristics. The most noteworthy challenge stems from Strawson's contention in Chapter 2 of *Individuals* (1959) that auditory experience, in contrast to visual and tactile-kinaesthetic experience, is not intrinsically spatial. Strawson concludes that a purely auditory experience would not include the experience of space. I evaluate several ways to understand this claim, from the extreme view of O'Shaughnessy (2000) to Strawson's own more moderate view, and find that each fails. I argue that Strawson's thought experiments fail to establish that audition differs from vision and tactile-kinaesthetic experience when it comes to the intrinsic capacity to experience space.

Room therefore exists for a view according to which audition is a robustly spatial perceptual modality. This proposal explains the experiential ground of beliefs about the locations of things and

events and the spatial characteristics of environments that are formed on the strength of hearing.¹ In particular, that ground includes the auditory experience of sounds as enjoying relatively stable distal locations in the neighborhood of their sources.

My phenomenological claim does not imply that sounds are where we experience them to be. Though phenomenology imposes conditions on its accuracy, auditory experience might be mistaken when it comes to representing the spatial locations of sounds. The main thesis of this paper does, however, imply that any view according to which locational hearing is not systematically illusory faces difficult questions about the compatibility of the commonplace view that sounds are traveling compression waves with the facts concerning how auditory experience represents sounds in space.

2. Spatial hearing

My suggestion is that learning by means of audition about the locations of things and events in one's environment depends upon the spatiality of audition. That is, audition's capacity to represent space grounds auditory learning of spatial locations and relationships. Put simply, locational learning sometimes amounts to the formation of auditory perceptual beliefs because audition itself represents space.

Auditory perception of spatial characteristics is neither as accurate (it is more often off target) nor as precise (its targeting is not as fine-grained) as vision, but audition does involve the acquisition of information about the relative locations of audible events and objects. Spatial hearing researcher Jens Blauert says:

Research has shown that the region of most precise spatial hearing lies in, or close to, the forward direction and that, within this region, a lateral displacement of the sound source most easily leads to a change in the position of the auditory event.... The spatial resolution limit of the auditory system [about 1 degree of arc] is, then, about two orders of magnitude less than that of the visual system, which is capable of distinguishing changes of angle of less than one minute of arc (Blauert 1997, pp. 38-9).

But the spatial characteristics of audition are not limited to directional information. Concerning what he calls "distance hearing," Blauert reports, "For familiar signals such as human speech at its normal loudness, the distance of the auditory event corresponds quite well to that of the sound source" (1997, pp.

¹ Though I am most concerned specifically with the spatial and locational aspects of such beliefs, formed on the strength of audition, the explanation stemming from my proposal need not be exhaustive. Further conditions concerning, for example, the role of auditory recognition and identification of sound sources, might be required. Thanks to an anonymous referee for pushing this point.

45-6). Blauert notes that even though distance localization is much less accurate for unfamiliar sounds, including "unusual types of speech," even in such cases, "The auditory event is, to be sure, precisely spatially located" (pp. 45-6). Sound engineers make use of facts about localization to shape the experience of sounds in concert halls, movie theaters, and living rooms (see, e.g., Shinn-Cunningham 2003). Even a recent police tip sheet entitled "How to Be a Good Witness" instructs individuals to "Look in the direction of the sound--make a mental note of persons or vehicles in that area" (Kershaw 2002). This is representative of the intuitive and widely accepted view among auditory researchers that hearing informs subjects about the locations of things and events in egocentric space (canonical and representative examples of this view include, e.g., Shinn-Cunningham 2001a and 2001b, Bregman 1990, Hartmann and Wittenberg 1996, and the classic discussion in Mills 1972).

3. Experiencing the locations of sounds.

Hearing, I wish to argue, consciously furnishes information about the locations of objects and events in the environment *by* or *in* consciously presenting sounds themselves as located in the surrounding environment. You hear a car passing or hear a plate breaking behind you by hearing its sound to have some location in the surrounding space. The experience presents the sound as occurring at some distance and in a particular direction. You turn to look at the passing wreck or to see how much of a mess there is because, in a wide range of cases, sounds themselves seem to have more or less determinate locations outside the heads of their perceivers. This requires arguing that sounds are auditorily experienced as located and that hearing source locations depends upon hearing sound locations. This section and the next take up the former; Section 5 deals with the latter.

That ordinary hearing presents sounds not just as in some direction but also as at a distance is most easily illustrated by contrast with the case of listening with ordinary headphones.² Headphone listening differs from ordinary hearing in that sounds seem to come from somewhere between the subject's ears, and not from the environment. Stanley Gelfand describes headphone listening as involving what he calls *intracranial lateralization*: "Sounds presented through a pair of earphones are perceived to come from within the head, and their source appears to be lateralized along a plane between the two ears"

² Some headphones are capable of delivering a relatively realistic rendering of auditory space, including externalized sound images. Such headphones are often expensive, since the very best examples require customized fitting and signal processing that incorporates modeling of the wearer's pinnae. This is because signal transformations that depend upon reflections due to ear shape and head position are important cues to realistic spatial hearing. Less sophisticated designs can sometimes produce reasonable sound imaging, though this still requires powerful signal processing.

(1998, p. 374). Gelfand contrasts headphone listening with what he describes as *extracranial localization*: "Sounds heard in a sound field seem to be localized in the environment" (p. 374). Hartmann and Wittenberg claim, straightforwardly, "Listeners perceive the sounds of the real world to be externalized. The sound images are compact and correctly located in space" (1996, p. 3678).

My phenomenological claim is that, in a wide range of central cases, we experience a sound to be located in the neighborhood of its source.³ When we do not, as when a sound seems to "fill a room" or to "engulf" us, this does not show that audition is non-spatial or that the sound seems to lack location entirely. Rather, the sound auditorily appears to be "all around," or at least in a larger portion of the surrounding space. Even hearing a sound to be located "in the head" when listening to earphones is another sort of sound location experience, though a bit odd.

It is notoriously difficult to argue for a phenomenological claim. But that is what I hope to do because, as I have indicated, our account of audition's epistemic role and of its role in perceiving the world of things and events depend on this claim. If the phenomenological claim is correct, then audible objects and events are experienced to have audible locations because their sounds seem to have audible locations. One way to argue for the phenomenological claim, which I will deploy in the sections that follow, is to show that it emerges as the best among alternative descriptions of the phenomenology of locational hearing.

4. "Coming from"

Despite my phenomenological claim that sounds seem to perceivers to be located at some distance in a particular direction, it is natural to describe sounds as "coming from" their sources. We ask where the beeping sound is coming from and whether the hum is coming from above or below. If sounds seem to *come from* particular places, in a spatial sense of "coming from," then locatedness as I have characterized it does not accurately capture the phenomenology of spatial auditory experience. Sounds, in that case, do not seem to occupy relatively stable distal locations; they seem to come from particular locations.⁴

The virtue of this way of describing the phenomenology is that it is consistent with facts about our ability to acquire through auditory perception information about the locations of objects and events in the environment. If sounds seem to come from their sources, then we can learn about where those sources are through audition without hearing sounds themselves as located at their sources.

³ See also Pasnau 1999 and Casati and Dokic 2005 for similar claims.

⁴ By "stable" I mean that the sounds do not seem to change locations. One might substitute "stationary". This does not imply, for instance, that no other changes seem to take place at those locations.

How are we to take talk of sounds' being heard to "come from" a location? It might be that sounds are heard to come from a particular place by being heard first at that place, and then at successively closer intermediate locations. But this is not the case with ordinary hearing. Sounds are not heard to travel through the air as scientists have taught us that waves do. Imagine a scenario in which engineers have rigged a surround-sound speaker system to produce a sound that seems to be generated by a bell across the room. This sound subsequently seems to speed through the air toward you and to enter your head like an auditory missile. This would indeed be a strange experience, one unlike our ordinary experiences of sounds, which present them as stationary relative to the objects and events that are their sources.⁵

Perhaps, instead, sounds are heard to *be* nearby, but to have *come from* a particular place, much as a breeze seems to have come from a certain direction. But feeling a breeze is more like listening with an earphone. The tactile experience of feeling a breeze, like the auditory experience of hearing a sound played through one earphone, includes direction but not distance. Earphone listening differs from ordinary hearing in where sounds seem to come from because it differs in where sounds are heard to be. Imagine feeling *where the fan is* by feeling its breeze. Tactile spatial perception is not in this respect like ordinary spatial hearing.

Since sounds seem to come from sources in a sense that includes distance as well as direction, and not in a sense that includes travel, neither spatial understanding of sounds' seeming to come from their sources does justice to the phenomenology of ordinary auditory experience. Thus, neither spatial understanding of "coming from" explains how auditory perceptual experience actually furnishes information about the locations of sound sources. The best sense to make of sounds' seeming to come from particular locations is that they have *causal sources* in those locations--that they are *produced* or *generated* in those locations.

5. Spatial hearing without located sounds?

I have argued that sounds seem to have relatively stable distal locations, and that they do not seem to travel or come from their sources in any spatial sense. Part of the motivation for this line of argument is that our capacity to form beliefs about the locations of ordinary things and happenings in the environment on the basis of auditory experience requires explanation. If auditory experience has spatial content, if we

⁵ It is worth making clear that I do not believe that it is never possible to auditorily experience a sound as traveling. So, for instance, a sound might seem to travel if the train producing it travels. A sound might seem to travel toward you if the vehicle making it does. My claim is that sounds do not ordinarily seem to travel *relative to* or *independent from* their sources. Thus, a stationary source does not ordinarily give rise to a sound that seems to travel.

are auditorily aware of locations of sounds, and if sounds seem to be located near their sources, then the beginning of an explanation is in hand. What remains is to explain the mechanisms concerning how sounds are localized and how sounds are associated with sources. Empirical research on spatial hearing has provided answers to the former question based on the physics of sound waves and on our ability to determine direction and distance information primarily in virtue of differences in wave characteristics at the two ears. For instance, interaural time and level differences, head-related transfer functions, and secondary reflections all play a role in determining the spatial characteristics of one's auditory experience (see Blauert 1997).⁶

However, might such perceptual processes, which I have invoked as the mechanisms by which sounds are perceptually experienced to be located, serve instead to explain why ordinary objects and events auditorily seem to be located without adverting to the experience of located sounds?

One promising approach based on this idea accepts that hearing is a spatial experience, but rejects the phenomenological claim as I have offered it. The approach is to say that we hear sounds to have pitch, timbre, and loudness, though not as having location. Rather, we hear *ordinary events and objects as located* and as the generators or sources of audible qualities that lack spatial characteristics entirely. We simply fail to perceive the locations of sounds. Sounds on this view are not heard to have locations, they are heard to have located sources. This claim is supplemented by an explanation in terms of sound wave transmission for how we localize sources in the environment. A location is thereby ascribed to a source, though not to a sound.

This description provides an account of the spatial phenomenology of auditory experience consistent with our ability to form auditory beliefs about source locations without claiming that sounds are heard to have locations. To see why it fails we need to consider the way in which spatial auditory experience conveys perceptually available information about the locations where sounds are generated.

I have claimed that hearing provides information about the ordinary objects and events that surround us--notably, information about where those things are and occur. However, sounds are, in a

⁶ I will not here address the latter set of questions, concerning how we associate sounds with sources. I treat these issues in O'Callaghan (2007, Chapter 11) and in forthcoming work. Location may be an important component to an account of how sounds are associated with sources. But, the auditory system is functionally segregated into "what" and "where" pathways analogous to those of vision (see, e.g., Kubovy and Van Valkenburg 2001), so it is reasonable to suppose that locational and identificatory processes comprise distinct if complementary roles when it comes to associating sounds with sources. My concern here is with the contribution of processes and information concerning the locations of sounds and sources, and not with an exhaustive explanation of that which grounds recognition and identification.

relatively innocuous sense, the immediate objects of auditory experience. Whatever else one hears, such as a bell, a collision, or a speech, one hears it in virtue of hearing a sound. So, though we sometimes hear things other than sounds, we hear those things in virtue of hearing some sound. Experience of a sound consists in the experience of temporally extended complexes of audible qualities including pitch, timbre, and loudness. So, there is an important sense in which whatever you hear, you hear it in virtue of enjoying an experience of audible qualities such as pitch, timbre, and loudness.⁷

The account we are considering is that we hear objects and events as located by means of the sounds they generate, though we do not hear those sounds as located. One's experience of non-located pitch, timbre, and loudness must therefore ground one's auditory experience of the locations of ordinary objects and events.

I claim, however, that we cannot *hear* just non-located audible qualities and located objects, *full stop*. This would amount to a precarious perceptual situation. How could hearing as of non-located qualities furnish perceptually manifested locational information about sound sources?

One way already has been mentioned. Locational information might be encoded temporally, for example, in terms of time delays between waves reaching the ears. However, since the suggestion is that we are auditorily *aware* of the locations of things and happenings, this information must be conveyed somehow in conscious perceptual experience. That is, the ordinary objects and events must be represented in auditory experience as located, and this must be accomplished by means of awareness of audible qualities. Suppose, for illustration, that the fundamental audible qualities comprise pitch, timbre, and loudness, so that an auditory experience conveys information about the locations of material objects and events by means of one's awareness of complexes of these basic attributes. Temporally encoded location information, then, must be manifested through one's experience of pitch, timbre, and loudness. For a spatial experience grounded in the audible qualities to be an auditory experience of location, the audible qualities must themselves bear spatial information. Since, as I have argued, sounds and their audible qualities do not auditorily seem to come from particular locations in a sense that involves travel or arrival, auditory awareness of location must occur thanks to an auditory awareness of located audible qualities. Sounds, the bearers of audible qualities, must appear to occupy stable distal locations if we are to learn of those locations through spatial auditory experience.⁸

⁷ One might, of course, hear white noise or other aperiodic sounds that lack determinate pitch. Such sounds nonetheless have loudness and distinctive timbre.

⁸ One might attempt to press into service the distinction between the "what" and "where" auditory pathways in defense of the initial suggestion. Suppose the "what" system identifies the sound and its source on the basis of audible qualities stripped of locational significance. A consciously accessible

Another way to put this point is as follows. Some feature of the sound in virtue of which one hears a thing or event must ground one's experience of its spatial location. This is because, in general, one cannot experience some particular as being at a place unless one experiences some of its (audible, visible, perceptible) qualities as being located at that place. The audible qualities in virtue of which one hears things and events are temporally extended combinations of pitch, timbre, and loudness. Since the experience of non-located audible qualities such as pitch, timbre, and loudness cannot ground a spatial auditory experience as of located things and events, the experience of located audible qualities must do so. But the experience as of located pitch-timbre-loudness complexes just is the experience as of a located sound, since sounds are the primary bearers of pitch, timbre and loudness. The problem with the proposal being considered is that it fails to provide a coherent account of spatial hearing and the mediatedness of object and event awareness in audition.⁹

A distinction thus can be drawn between hearing sounds themselves as located and perceiving information about the locations of material objects, stuffs, and events in the environment by means of audition. Given that we learn the locations of ordinary objects and events in audition, the question is whether the latter would be possible without the former. Since sounds are the immediate objects of auditory experience, since auditory awareness of location occurs by means of awareness of audible qualities, and since sounds seem to come from their sources only in a causal sense, hearing sounds and their qualities as located is required in order to experientially represent the locations of material objects and events through audition. Sounds are heard to have locations, by means of which they provide perceptual information about the locations of their sources.

6. Skepticism about spatial audition

auditory representation of the audible source might then simply be ascribed a location without adverting to the representation of a located sound. But, in what does the consciously accessible representation of an audible sound source as at a location amount to if it does not include a bound collection of audible qualities as at that location? One might escape the dilemma at the considerable expense of denying either that we ever hear sounds or that we hear things by hearing their sounds, so that one might hear a thing without hearing a sound. But that is not the suggestion we are considering, and it is sufficiently revisionary as to demand independent motivation and evaluation. I am comfortable to set it aside.

⁹ To be clear, I am not denying that sounds themselves are, metaphysically, some variety of object or event. In fact, I, like Casati and Dokic (2005), believe sounds are distally located events (see O'Callaghan 2007, forthcoming 2008). What I mean is that awareness of objects and events that are not themselves sounds, such as ordinary bells and accidents, depends upon one's awareness of a sound.

In the first half of this paper, I have developed and defended a view according to which learning on the basis of audition about the locations of things and events in one's environment takes place thanks to the spatial character of auditory experience. In particular, one's experience of sounds as distally located and stationary relative to their apparent sources grounds the formation of beliefs about the arrangement of the sound-producing "furniture".

Philosophers, however, have been skeptical about the spatial character of audition. Perhaps the most famous skeptic about audition's capacity to provide a grip on space is Strawson (1959). In "Sounds", Chapter 2 of *Individuals*, Strawson concludes that a creature who enjoyed only auditory experiences would lack the means for spatial experience.

...It does not follow that where experience is supposed to be exclusively auditory in character, there would be any place for spatial concepts at all. I think it is obvious there would be no such place. The only objects of sense-experience would be sounds. Sounds of course have temporal relations to each other, and may vary in certain ways: in loudness, pitch, and timbre. But they have no intrinsic spatial characteristics: such expressions as "to the left of", "spatially above", "nearer", "farther" have no intrinsic auditory significance.... A purely auditory concept of space... is an impossibility (pp. 65-6).

Strawson is clear to contrast audition with vision and tactile-kinaesthetic experience. The latter furnish the materials for spatial concepts because visual and tactile-kinaesthetic experiences are intrinsically spatial--one cannot enjoy a visual or a tactile-kinaesthetic experience without experiencing spatial relationships. Audition, according to Strawson, has no such intrinsic spatial character.

Though Strawson gives little argument for this claim, many have been convinced. Matthew Nudds echoes and endorses Strawson's remarks.

In characterizing the contrast between auditory and visual experience in the way that I have I am attempting to draw attention not merely to the fact that the auditory experience has a different spatial structure, but that it has a structure that is non-spatial. We are visually aware of the space within which we see things in a way that we are not auditorily aware of the space in which we hear sounds. When we see (or seem to see) something, we see it as occupying or as located within a region of space; when we hear (or appear to hear) a sound we simply hear the sound, and we don't experience it as standing in any relation to the space it may in fact occupy (Nudds 2001, pp. 213-4).

So does Brian O'Shaughnessy: "...We absolutely never immediately perceive sounds to be *at* any place. (Inference from auditory data being another thing.)" (2000, p. 446).

These authors claim that audition has no spatial content entirely its own; in isolation, it is constitutively incapable of presenting spatial relationships. Consequently, audition furnishes information

about the spatial features of an environment only in virtue of one's prior grasp on spatial properties and relations gleaned from vision or tactile-kinaesthetic experience. Audition informs us about space, if at all, only thanks to correlation or correspondence with spatial experience in other modalities.

The point of this paper is that Strawson and the others are wrong about audition.

It is natural to believe, as I mentioned at the outset, that hearing frequently informs us about the locations of things and events in our environment. We turn to look toward a familiar voice or a creaking door because audition bears information about *where* those activities take place around us. On the strength of audition we form beliefs about the locations of items and occurrences in our surroundings. I have argued not only that audition furnishes subjects with information about the locations of things and events, but also that we learn about such locations *by* or *in* hearing their sounds to be at some distance and in a particular direction--that is, by hearing sounds as having locations. I have argued that audition is a thoroughly spatial perceptual modality.

Skepticism about audition's capacity to represent sounds in space may take a number of forms, but each depends on the idea that auditory experience is not intrinsically spatial or that spatial concepts have no intrinsically auditory significance. But what does it mean to say that auditory experience or its objects are not intrinsically spatial? The most extreme form of the claim, represented by the O'Shaughnessy quote above, implies that no auditory experience represents spatial location because audition lacks the capacity to represent space altogether.

7. Non-spatial hearing?

For O'Shaughnessy (2000), to say that we hear the locations of things or the places sounds come from is a mistake. We never strictly hear or auditorily experience such locations. At best, we associate auditory features with extra-auditory experiences of location--experiences in other modalities that possess genuine, non-derivative spatial content--and work out where sounds come from by inference from "auditory data". O'Shaughnessy's answer to the question that motivates this paper is that we become informed of locations through audition thanks to a cognitive or inferential process based on non-spatial "auditory data".¹⁰

¹⁰ In fact, Strawson agrees with part of this. Strawson says that assigning location on the basis of audition is "sufficiently explained by the existence of correlations between the variations of which sound is intrinsically capable and other non-auditory features of our sense-experience" (p. 66). But Strawson does not think this is a matter of inference or facts accessible to the subject, as O'Shaughnessy seems to think. Strawson says,

I do not mean that we first note these correlations and then make inductive inferences on the basis of such observation; nor even that we could on reflection give them as reasons for the

On this view, audition helps us to work out the locations of things and events we hear because audible features of sounds provide clues to spatial features and locations revealed through other sense modalities. The pressing question, of course, is which features of sounds provide reliable inference cues about locations perceived only through other senses. How does non-spatial auditory experience so reliably furnish the materials for accurate beliefs concerning the spatial features of one's environment?

Notice that this is not the view considered in Section 5, according to which audition is spatial without presenting *sounds* as located. Instead, we are now asking whether audition and the awareness of sounds might provide information about location without presenting spatial characteristics. On this view, qualities such as pitch, timbre, and loudness provide clues to the egocentric locations of their sources without seeming located.

Strawson, however, is right about this much: pitch and timbre *in themselves* provide little useful information about location. A flute may auditorily appear to have the same pitch and timbre whether it is right in front of me or across the street. A tuba's peculiar pitches as such give no clue about its location. Even Doppler-shifted pitches offer little help since a stationary source might mimic them.

Surely loudness is a different story, since nearer things make louder sounds than things at a distance? Loudness, or changes to loudness, may indeed provide useful clues about location, but there are no strict rules governing the relationship between loudness and distance. A quiet tuba may be playing nearby and a loud tuba may be far off. In fact, given loudness constancy effects, a tuba playing at constant volume may seem to have the same loudness when playing nearby or at a distance (see Zahorik and Wightman 2001). Loudness provides less useful information about distance than one might hope.

assignments of distance and direction that we in fact make on the strength of hearing alone (p. 66).

Still, Strawson claims that "the *de facto* existence of such correlations is a necessary condition of our assigning distances and directions on the strength of hearing alone," (p. 66) and so distinguishes hearing from vision. Evans (1982) also rejects the inference view. He says, "When we hear a sound as coming from a certain direction, we do not have to *think* or *calculate* which way to turn our heads (say) in order to look for the source of the sound" (p. 155). But he does hold that we might "envisage an organism which can be conditioned to respond differentially to those different values of the proximal stimulus which code the direction of sound, for instance by pressing a button" (pp. 154-5). Auditory spatial content requires using such correlations to direct spatial behavior. Nudds (2001, p. 226, note 3) suggests his arguments are compatible with such a view. To the extent that these accounts require extracting locational information from non-spatial audible qualities of sounds whose changes correspond to changes in location, my arguments against O'Shaughnessy apply. My blindsight argument, below in the text, does not.

Loudness, pitch, and timbre are even less useful where audition is best. None of these audible qualities is capable, in itself, of informing listeners which direction to turn to find an oncoming vehicle. Learning about location on the basis of audition simply cannot take place as a result of something gleaned just from information about the audible qualities of sounds. The dimensions on which pitch, timbre, and loudness vary are not ones to which changes in location correspond.

It is difficult to imagine what audible feature of a sound could provide information about the location of a source without that feature's seeming to be located. No perceptible qualitative attributes, such as pitch, timbre, or loudness, correspond reliably to egocentric location, and variations in qualities do not correspond reliably to changes in egocentric location. Pitch and timbre are independent of spatial location, and loudness is independent of direction and probably distance.

Might some other attribute of sounds, which is not consciously accessible, serve as the auditory basis for learning about locations? Features that are not consciously experienced but which ground the mechanisms of localization *can* furnish location information. For instance, temporal information including wave onset delays and other properties of sound waves might furnish the information required for audition-based spatial beliefs without invoking the auditory experience of space.

We are approaching an impasse. In my view, this strategy simply fails to capture the empirical and phenomenological facts. Location information detected by auditory processes is *perceptually* and *consciously* manifested. It is not merely delivered subconsciously to epistemic processes that manage belief formation. Though the mechanisms of auditory localization utilize temporal information, such as wave onset delays, and phase information, subconscious localization processes result in conscious perceptual access to locations. This is the *experience of audible* direction and distance from oneself.

Support for this claim goes beyond introspection. Behavioral evidence from performance shows clearly that localization is a low-level, automatic perceptual process that does not require inference, effort, or extra-auditory cognition to "work out". Turning reactively to a sound behind the head is virtually unavoidable. Furthermore, the point is straightforward when we consider a contrast with the case of blindsight (see, for instance, Weiskrantz 1986, Block 1995, or Boyer, Harrison, and Ro 2005). Blindsighted subjects reliably report beliefs or guesses about, for example, the location of a flashcard without reporting visually experiencing the location of the flashcard. Suppose that auditory experience is non-spatial and that we do not hear sounds as located in any way. We instead acquire information about the locations of sound sources through an entirely unconscious process something like what occurs in blindsight. Location information is thus represented through a course of subconscious perception. In this case, audible qualities do not bear location information, since by hypothesis locational information is not experientially manifested. Subjects simply find themselves believing that sound sources have certain locations without reflective awareness of experiencing the locatedness of those sources. This story is

implausible as a characterization of ordinary hearing's phenomenology. Subjects who report and act upon the audible locations of events in egocentric space neither guess nor report beliefs lacking ground in auditory experience. It *seems* we hear where things are and where sounds come from. Ordinary hearing differs from the auditory analog of blindsight in that audible qualities seem to have audible locations--a fact demonstrated by behavioral and introspective tasks. If this appearance is illusory and the blindsight explanation is plausible, then we have a fresh example of reliable perceptual beliefs formed in the absence of genuine conscious awareness. This itself seems improbable. It is far less improbable that we are simply *aware* of locations in audition.

This extreme form of skepticism about spatial audition, according to which audition is entirely non-spatial, therefore fails. How else might we understand the claim that audition is not intrinsically spatial?

8. Two conceptions of the pure or exclusively auditory experience

When Strawson claims that auditory experience is not intrinsically spatial, he does not wish to rule out hearing the locations of sounds. He does, after all, say that ordinary perceivers assign directions and distances to sounds "on the strength of hearing alone" and that he concedes the "full force" of this phrase (p. 66). Under ordinary conditions, audition does have spatial content. The claim, rather, is that audition's locational content is somehow parasitic upon the spatial content of another modality; audition represents space only thanks to spatial experience in some other modality. Audition, on this picture, comes to have spatial content, but its coming to have that content requires extra-auditory resources. Experience with another modality is a condition necessary for the experience of space through audition. Vision and tactile-kinaesthetic perception, however, require no outside resources.

Support for audition's spatial dependence stems from a thought experiment involving a purely or exclusively auditory experience. Strawson famously claims that a purely auditory experience, unlike a purely visual or purely tactile-kinaesthetic experience, would be a non-spatial experience. The world of sounds would be a No-Space world.

But Strawson's phrases "purely auditory experience" and "exclusively auditory experience" are vague. What constitutes such an experience? I will suggest that there are two possible ways to understand these phrases. However, neither understanding grounds an argument that yields Strawson's conclusion that the purely or exclusively auditory experience differs from an analogous visual or tactile-kinaesthetic experience in being non-spatial.

The most natural way to understand Strawson's notion of the exclusively auditory experience is that the subject of such an experience has only experiences that count as auditory, and no experiences that count as visual, tactile, kinaesthetic, olfactory, or gustatory. The purely auditory experience, according to

this understanding, consists exclusively in the experience of sounds throughout the subject's entire existence.¹¹

Would such an experience contain the materials for a conception of space? Despite Strawson's assertions, the question is empirical. And the empirical evidence points to "yes". Research into locational hearing over the past few decades demonstrates at minimum that audition has a prominent directional component that depends upon, for instance, differences in arrival time, phase, and level of sound waves at the two ears. In addition, hearing in the presence of sufficient binaural cues involves the "externalization" of sounds--sounds seem to come from somewhere outside the head. Headphone listening, in which sounds seem "internalized", or to come from somewhere inside the head, provides the useful contrast.

It is plausible to think that a subject's capacity to enjoy the rich spatial experience of ordinary hearing depends on having the right sort of visual or tactile-kinaesthetic experiences, or on having the capacity to interact with one's environment and to grasp the ways in which sensory stimulation would vary depending on one's movement (see Noë 2004). But it is also plausible to think that even the rich spatial experience of ordinary vision requires more than just visual stimulation. The capacity to move and to interact with one's environment, and the tactile-kinaesthetic experience that accompanies such movement surely makes visual experience richer. And vision surely enhances even tactile spatial experience of, for example, the substantial volumes of objects.

But there is very little reason to think that an exclusively auditory experience would leave absolutely no room to experience spatial characteristics, though vision would leave room. Imagine a subject in the following very simple situation. Wearing a pair of headphones but privy to no other experiences (in particular, neither visual, tactual, nor kinaesthetic), our subject is presented with a brief tone through one headphone, followed by a brief tone through the other. It is difficult to imagine such an experience failing to provide even the minimal experience of space. Introducing a fade from one headphone to the other should result in the experience of movement, along what we would describe as a plane between the ears, despite the subject's inability to express it just that way. I can see little reason to accept that, when exposed auditorily to the blooming, buzzing world, the same subject's experience, though impoverished in many respects, would lack even the most minimal spatial character. At the very least, this is not a matter to be decided by purely conceptual considerations. Though difficult to test, basic

¹¹ I have been speaking as if we do have auditory access to ordinary things and events, which, according to Strawson, requires the perception of space (which is required to ground the conception of existence unperceived necessary for experiencing objective particulars). Strawson correctly points out that auditory awareness of ordinary things and events is mediated by one's awareness of a sound. Whatever else we hear, such as cars and collisions, we hear it by way of hearing a sound.

facts about the physiology and psychology of hearing suggest that the entirely auditory experience need not be non-spatial.

Perhaps what Strawson means by the purely or exclusively auditory experience is one that counts as what we might call a "minimal auditory experience". That is, such an experience meets the minimum requirements on an experience's being an auditory experience. For instance, awareness of a uniform sinusoidal tone at a given loudness with none of the cues required for binaural spatialization might count as minimally auditory without furnishing the materials for concepts of spatial relationships. The idea is that such an experience, though auditory, lacks some aspects or dimensions of ordinary auditory experience, or of the entirely auditory experience discussed above. Such a minimal auditory experience perhaps lacks spatial characteristics for reasons Strawson cites. Pitch and other audible qualities are not obviously intrinsically spatial in the way that, for example, shapes are. They can at least be superficially characterized without mention of spatial notions. But the independence of concepts of audition's objects from concepts of spatial relationships does not guarantee both *that a purely auditory experience would be non-spatial* and *that a purely visual or tactile-kinaesthetic experience would not be non-spatial*.

It is unclear whether it is correct to describe such a minimal auditory experience as non-spatial, or whether it ought instead to be described as the experience of sound "all around". Though the subject might fail to extract spatial concepts from such an experience, that does not entail a non-spatial experience.

Fortunately, the choice does not matter. For we can imagine a "minimal visual experience" in which the same lack of clarity applies. Imagine a subject who enjoys no non-visual experiences, but who is made to experience a uniform-grey ganzfeld as a result of the application of halved table-tennis balls over the eyes. (Or imagine viewing just a very bright, uniform flash of light.) Such homogeneous visual awareness of light may not furnish its subject with the materials for distinguishing places or spatial relationships. But, though visual, it is not clear whether this experience should be described as non-spatial or as the experience of uniform grey all around. Likewise, we might construct a similar "minimal tactile experience" in which the uniform feeling of warmth on every skin surface fails to ground spatial concepts, and to which it is unclear whether or not to ascribe a spatial character. There is nothing special about the auditory case.

So, understanding the purely auditory experience as what I have called a "minimal auditory experience" perhaps shows that such an experience would have a non-spatial character, and that it would be insufficient for grounding spatial concepts. But, the possibility of analogous "minimal visual experiences" and "minimal tactile experiences" shows roughly the same thing for purely visual and purely tactile experiences. This understanding of the purely auditory experience cannot be deployed in service of Strawson's assertion.

Strawson's considerations therefore do not show that a purely or exclusively auditory experience would be non-spatial, and that audition differs from vision and tactile-kinaesthetic experience in this respect. They do not show this because there are two ways to understand what a purely or exclusively auditory experience involves. On the first, the subject enjoys no experiences other than auditory experiences. On the second, the subject enjoys an experience that counts as an auditory experience, but which may be severely impoverished. If the first, empirical evidence suggests the subject could have a spatial experience. If the second, analogous impoverished visual and tactile-kinaesthetic experiences (for example, of light and heat, respectively) can be conceived which also are non-spatial.¹² Strawson's conclusion trades on the ambiguity.

Therefore, the claim that a purely or exclusively auditory experience--an experience of sounds alone--would be entirely non-spatial, and the claim that this constitutes a significant respect in which audition differs from vision and tactile-kinaesthetic experience, rests on a mistaken understanding of audition.

9. The structure of auditory space

It is long-standing tradition to model vision in spatial terms. Visual experience has been likened to pictures, photographs, and movies. The telephone and phonograph, on the other hand, are thought to capture what is most interesting about auditory experience. It is indeed difficult to characterize the most striking features of vision in non-spatial terms, and it is tempting to think the telephone shows that the same does not hold for audition. Surely this intuitive difference between vision and audition, which is about the role of space, moved Strawson to distinguish hearing from visual and tactile-kinaesthetic experience and to say that audition is not an intrinsically spatial perceptual modality.

Nudds (2001) attempts to explicate Strawson's contention that audition is not intrinsically spatial in terms of the claim that vision involves a spatial "field" but audition does not.¹³ That is, vision has, and audition lacks, a spatial *structure*.

¹² One might suggest that this is because, though they can be experienced through visual or tactile-kinaesthetic means, neither light nor heat has intrinsic spatial characteristics. This, as in the case of audible qualities such as pitch, is implausible if light and heat, for example, are metaphysically dependent upon physical attributes such as energy and frequency, since such physical attributes arguably are intrinsically spatial.

¹³ See Clark (1996, 2000) for very helpful discussions of the notion of a "visual field". Clark highlights the revisionary import of various ways of explicating the spatial aspects of visual experience in terms of a visual field.

The argument goes like this. We know vision involves a spatial field because the field can be experienced independent of the objects that occupy it--one can experience places *as places* independent of experiencing objects to occupy them. Possessing spatial structure indeed depends, for Nudds, on the capacity to represent empty places. Nudds argues that audition lacks this capacity. In vision, he claims, there is an experiential difference between enjoying awareness as of a place that might be, but is not, occupied by an object and failing to enjoy awareness of anything to occupy a place where something might have been experienced. The difference is between *experiencing as empty* and *not experiencing*.¹⁴ Drawing on an example of Martin's (1992, p. 199), he distinguishes the phenomenon of experiencing a place as a place where something could be--as when one sees the hole in a donut as a place where something could be seen--from failing to experience something at a given place. According to Nudds, though this distinction applies to visual experience, audition makes no room for such a distinction: "There's no difference between not experiencing a sound at some place, and experiencing no sound there" (p. 213). We therefore are incapable of auditory awareness of empty places as such. Audition is not intrinsically spatial since it lacks the spatial structure evinced by awareness of empty places.

In what does the distinction consist between experiencing a place as an empty place where something could be and not experiencing anything at that same place? What is the difference between not experiencing anything in the middle of the donut and experiencing the middle as an empty place? The difference, if anything, is a difference in ways of attending. One can attend to the spot inside the ring as a place that might be occupied, or one can fail to attend to it. When you fail to attend to it, you do not experience it as a place that could be occupied. You simply fail to experience anything there at all. So, what this view amounts to is that possessing spatial structure is a matter of possessing the capacity to attend to a place without attending to some object that occupies it. If subjects could attend only to places as occupied by objects, the distinction would not take hold.

But subjects do possess the capacity to auditorily attend to places a sound could be but is not. I might attend to a place as the source of no sound, and this is different from failing to experience a sound at that place while attending to sounds elsewhere. Suppose you hear music from a single speaker. You might auditorily attend to the sound's location. You might also reflect and recognize that the sound does not seem to be behind you. You might then auditorily attend to the space behind you and attend to it *as* a place where there are no audible qualities. You might then experience it as a place where a sound could be heard to be. Suppose a second speaker behind you begins to play. The place you formerly experienced

¹⁴ This can be seen as a difference in quantifier scope when characterizing the experiences. In particular, it is a matter of whether the proposition that encodes the content of the experience includes the negative existential or whether the negated existential has widest scope in a proposition describing the experience.

as a soundless place that nonetheless could be occupied by sound you now experience to contain a sound. In fact, auditory attention to a place enhances one's ability to detect features of sounds that subsequently occur at that place. Auditory attention is agile and encompasses a field whose spatial structure rivals vision's.

It would, nevertheless, be foolish to deny that spatial experience in vision differs from spatial auditory experience. But the differences do not show that the *structure* of auditory space differs radically from that of vision, or that audition lacks spatial structure entirely. One difference is that, ordinarily, visual features or objects can at any time be detected in all directions (though not at all distances) within the sweep of one's forward gaze (see Matthen 2004). Frequently, however, many auditory directions are devoid of sounds and audible features, though audition affords awareness of a much more encompassing space than vision, including things behind one's head. We would find it difficult even to characterize this difference without assuming spatial auditory structure. The grain of truth in the idea that audition lacks intrinsic spatial structure, granting differences in acuity and populousness relative to vision, is that we do not in ordinary circumstances experience sounds to exhibit *internal* spatial characteristics. That is, sounds, unlike material objects, do not seem to have determinate shapes or sizes, though they may seem to occupy regions of space. This failure to experience in the objects of audition intra-object spatial features such as three-dimensional shape does not imply that auditory experience is not itself intrinsically spatial or that audition lacks the capacity to represent space.¹⁵ Though features of the peripheral auditory system (which has only two eardrums, and not a pair of two-dimensional retinal arrays) limit the *internal* spatial properties sounds themselves are perceived to have, a complex set of processing strategies makes possible the representation of sounds as having *locations* in three-dimensional space. Audition represents sounds to occur at locations characterized by distance and direction in egocentric space. Thus we hear the world of sounds to be a world with space.

I have argued against the view that audition is spatial only parasitically or in a way that depends upon the spatial content of some other modality; I have argued against the view that audition is not intrinsically spatially; and I have argued against the view that audition is simply non-spatial. This exhausts the obvious candidates for skepticism about spatial audition. Nothing conceptual shows that audition lacks the materials necessary for the representation of space. Audition, I conclude, enjoys rich spatial content.

¹⁵ It may, however, imply that spatial features, such as spatial boundaries, are less critical to the individuation of audition's objects than to the individuation of visual objects. See Kubovy and Van Valkenburg 2001.

10. Conclusion

One of the primary functions of audition is to inform us about things and happenings in the world. Among these are the very things and happenings about which vision informs us. Furnishing locational information is a central part of performing this function, not least for its utility in navigating one's surroundings when vision is otherwise occupied. Audition, like vision, provides access to features and particulars which populate a three-dimensional spatial field. Though audition and vision inform us about different aspects of objects and events, they inform us in recognizably complementary ways about the very same world.

We learn about the locations of things and events in our environments through hearing because we hear their sounds to be located. Not only is auditory experience itself spatial, but in audition sounds and audible qualities themselves are experienced to have spatial locations. This involves a claim about the phenomenology of audition: that sounds under a wide variety of conditions seem to us to be located at a distance and in a given direction. I have used this phenomenological claim to ground a particular proposal about just how auditory experience furnishes information about the locations of ordinary things and events in the surrounding environment--it furnishes such information in virtue of presenting sounds as having relatively stable distal locations. So, I have made three proposals. First, audition furnishes subjects with information about the locations of things and events because auditory experience itself is spatial. Second, we hear the locations of things and events by or in hearing locational information about sounds. Third, we auditorily experience sounds themselves as having relatively stable distal locations.¹⁶

References

Blauert, J. 1997. *Spatial Hearing: The Psychophysics of Human Sound Localization*. Cambridge, MA: MIT Press.

Block, N. 1995. On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18: 227-47.

Boyer, J. L., Harrison, S., and Ro, T. 2005. Unconscious processing of orientation and color without primary visual cortex. *Proceedings of the National Academy of Science*, 102(46): 1675-9.

¹⁶ Thanks to Barbara Shinn-Cunningham for several enjoyable conversations about spatial audition and the auditory experience of space. I would also like to thank two anonymous referees for extensive and extremely useful comments.

- Bregman, A. S. 1990. *Auditory Scene Analysis: The Perceptual Organization of Sound*. Cambridge, MA: MIT Press.
- Casati, R. and Dokic, J. 2005. Sounds. In Zalta, E. N. (ed.), *The Stanford Encyclopedia of Philosophy*.
- Clark, A. 1996. Three varieties of visual field. *Philosophical Psychology*, 9(4): 477-95.
- Clark, A. 2000. *A Theory of Sentience*. Oxford: Oxford University Press.
- Evans, G. 1982. *The Varieties of Reference*. Oxford: Oxford University Press.
- Gelfand, S. A. 1998. *Hearing: An Introduction to Psychological and Physiological Acoustics*. 3rd edition. New York: Marcel Dekker.
- Hartmann, W. A. and Wittenberg, A. 1996. On the externalization of sound images. *Journal of the Acoustical Society of America*, 99(6): 3678-88.
- Kershaw, S. 2002. Feeling that witnesses need a hand, police offer one. *The New York Times*, October 16, 2002.
- Kubovy, M. and Van Valkenburg, D. 2001. Auditory and visual objects. *Cognition*, 80: 97-126.
- Martin, M. G. F. 1992. Sight and touch. In Crane, T. (ed.), *The Contents of Experience*. Cambridge: Cambridge University Press.
- Matthen, M. 2004. Features, places, and things: reflections on Austen Clark's theory of sentience. *Philosophical Psychology*, 17(4): 497-518.
- Mills, A. W. 1972. Auditory localization. In Tobias, J. V. (ed.), *Foundations of Modern Auditory Theory*, volume II, 303-48. New York: Academic Press.
- Noë, A. 2004. *Action in Perception*. Cambridge, MA: MIT Press.
- Nudds, M. 2001. Experiencing the production of sounds. *European Journal of Philosophy*, 9: 210-29.

- O'Callaghan, C. 2007. *Sounds: A Philosophical Theory*. Oxford: Oxford University Press.
- O'Callaghan, C. forthcoming 2008. Constructing a theory of sounds. *Oxford Studies in Metaphysics*, 4.
- O'Shaughnessy, B. 2000. *Consciousness and the World*. Oxford: Oxford University Press.
- Pasnau, R. 1999. What is sound? *Philosophical Quarterly*, 49: 309-24.
- Shinn-Cunningham, B. 2001a. Creating three dimensions in virtual auditory displays. In Smith, M. J., Salvendy, G., Harris, D., and Koubek, R. J. (eds.), *Usability Evaluation and Interface Design: Cognitive Engineering, Intelligent Agents and Virtual Reality*, 604-8. New Jersey: Lawrence Erlbaum.
- Shinn-Cunningham, B. 2001b. Localizing sound in rooms. In *Acoustic Rendering for Virtual Environments*, 17-22. ACM SIGGRAPH.
- Shinn-Cunningham, B. 2003. Acoustics and perception of sound in everyday environments. In *Proceedings of the 3rd International Workshop on Spatial Media*, 1-9.
- Strawson, P.F. (1959). *Individuals*. New York: Routledge.
- Weiskrantz, L. 1986. *Blindsight: A Case Study and Implications*. Oxford: Oxford University Press.
- Zahorik, P. and Wightman, F. 2001. Loudness constancy with varying sound source distance. *Nature Neuroscience*, 4: 78-83.