CONSTRUCTING A THEORY OF SOUNDS

Oxford Studies in Metaphysics, 5, 2009

Casey O'Callaghan

Rice University

Sounds and vision

Vision has dominated philosophical thinking about perceptual experience and the nature of its objects. Color has long been the focus of debates about the metaphysics of sensible qualities, and philosophers have struggled to articulate the conditions on the visual experience of mind-independent objects. With few notable exceptions, "visuocentrism" has shaped our understanding of the nature and functions of perception, and of our conception of its objects. The predominant line of thought from the early modern era to the present is that, in the philosophically interesting respects, as things are with vision, so they are with hearing, touch, olfaction, and the rest. A closely related line of thought has been particularly strong in the case of the secondary qualities. The more or less implicit assumption is that as things are with colors, so they are with sounds, tastes, and smells. This paper is predicated on skepticism about this kind of claim. I suggest that we put to rest the traditional lines of thought because hearing and the world of sounds are rich with raw material that presents both novel philosophical problems and telling new instances of old ones. The case of sounds and audition demonstrates that attention to modalities other than vision enriches our understanding of perception and the natures of its objects.

My goal in this paper is to develop the framework for a philosophical account of sounds. In particular, contrary to the traditional philosophical understanding of sounds as secondary qualities, and contrary to the commonplace scientific view that sounds are waves in a medium, I argue that sounds are *events* located in the environment near their sources. This proposal is designed, first and foremost, to capture the essentially temporal nature of sounds. Furthermore, it aims to explain the features of auditory perceptual experience in a way that avoids attributing widespread, systematic illusion. According to this account, sounds are particular individuals that bear audible qualities, persist, and travel only if their sources do. Perceivers hear publicly available, distally located sounds thanks to the waves that bear and transmit information about those sounds. I am a realist about sounds. Sounds are individuals in the world that possess many of the features we hear them to have. The proposal that sounds are events of this sort has

consequences for theorizing more broadly about perception. This theory of sounds discloses greater variety among the objects of perception than the traditional lines of thought imagine, and forces us to reconsider our visuocentric understanding of perception.

What kind of thing is a sound?

Sounds are public objects of auditory perception, I maintain until convinced otherwise. You might hallucinate a sound, but in that case you fail to hear a sound -- you just think that you do. In principle, others might hear any sound you hear. Tinnitus sufferers suffer hallucinations. Furthermore, I will assume that if you successfully hear anything at all, you hear a sound. Whatever else you hear, such as an object or a happening in your environment, you hear it by way of or in virtue of hearing the sounds it makes. Sounds are, in this innocuous sense, the immediate objects of auditory perception. This sense is innocuous because it is neutral on the question concerning whether you are, in another sense, immediately aware only of auditory sense data. Finally, sounds are frequently characterized by pitch, timbre, and loudness. This tells us very little about what kind of thing a sound is -- what ontological category it belongs to.

Sounds as properties

The traditional philosophical outlook has grouped sounds with the colors, tastes, smells, and other sensible attributes or secondary qualities. Popular analyses of such qualities then imply that sounds are either dispositions to cause auditory experiences in suitably equipped perceivers under the right sorts of circumstances; categorical bases of such dispositions; physical properties; simple, primitive, or manifest properties; or mere projections of qualities of experiences. The options are familiar from the literature on color.

Locke, for one, held that sounds are secondary qualities: powers, grounded in the primary qualities of bodies, to produce auditory experiences (1975: II, viii, 10). But to which *bodies* did Locke mean to attribute these powers? On a natural reading of the *Essay*, he meant to attribute them to sounding objects so that sounds, like colors, are dispositions ordinary objects have to affect perceivers' experiences (see, in particular, II, viii, 9-14). However, Locke may have spoken loosely and meant instead to attribute sounds to the medium that intervenes between object and perceiver so that sounds are dispositions of the medium itself, considered as a body, to

produce auditory experiences.[1] Depending which Locke meant, we get two views that differ on where sounds are located.

Robert Pasnau (1999) takes a stand on this issue concerning the locations of sounds. Pasnau introduces a view according to which sounds are properties of sounding objects, not of the medium. Sources themselves have or possess sounds on this view. For Pasnau, sounds either are identical with or supervene upon the vibrations of the objects we ordinarily count as sound sources, so sounds are *properties* that depend upon the categorical bases of Lockean powers. Pasnau and Locke thus both reflect the traditional understanding of sounds as secondary qualities or sensible attributes.

We can classify views developed in the spirit of the traditional model of sounds as sensible qualities according to their stance on two questions. (1) What is the correct account of the sensible qualities, in general? That is, are they dispositional properties, physical properties, or primitive properties with which perceptual experience acquaints us? (2) Are sounds properties of the medium or of the objects? A matrix of property views of sound results.

However, independent of providing the details of a philosophical account of sounds as sensible attributes, we need to ask whether this model is the right approach to sounds in the first place. I want in what follows to suggest that it is not. Both of the questions that yield the above matrix depend upon a misguided supposition.

The suggestion that sounds themselves are sensible qualities is attractive only if we are in a mood that overemphasizes similarities with color and entices us to provide an account that subsumes sounds with colors under a single metaphysical category. This should be resisted. Sounds themselves are not properties or qualities at all. Sounds are best understood as *particular*

---

[1] In a passage from the later *Elements of Natural Philosophy*, Locke (1823) says:

> That which is conveyed into the brain by the ear is called sound; though, in truth, till it come to reach and affect the perceptive part, it be nothing but motion. The motion, which produces in us the perception of sound, is a vibration of the air, caused by an exceeding short, but quick, tremulous motion of the body from which it is propagated; and therefore we consider and denominate them as bodies sounding.

individuals that possess the audible qualities of pitch, timbre, and loudness, perhaps along with other audible and inaudible properties. Sounds bear similarity and difference relations to each other that are based upon the complexes of audible qualities they instantiate. Sound sources, among which we count ordinary objects and events, such as bells, whistles, and collisions, make or produce sounds, but are not at intervals simply qualified by their sounds in the way that walls are qualified by colors.

Several kinds of consideration support this suggestion. First, sounds survive changes to their properties and qualities. A sound that begins high-pitched and loud may continue to exist though it changes to being low-pitched and soft. An object does not lose its sound and gain a new one when it goes from being high-pitched to being low-pitched, as with an emergency siren's wail. The sound of a spoken word begins with certain audible characteristics and ends with others, but a pitch shift is not the end of a sound. Determinate perceptible or sensible qualities, however, do not survive change in this way. The red color of the fence does not survive the whitewashing. The dank smell of the dog does not survive the perfuming. Particular individuals, such as the fence and the dog, however, survive changes to their qualities.

In addition, the identities of many recognizable sounds are tied to the pattern of audible qualities they exhibit over time. To be the sound of a duck's quack, or the sound of a spoken syllable, requires a certain complex pattern of changes in pitch, timbre, and loudness over time. The sound of the spoken word 'treatise' differs from the sound of the spoken word 'treason' precisely because each exhibits a different pattern of change in audible qualities over time.

Since sounds survive changes to their properties across time, sounds last through time. In particular, sounds have durations or lifetimes. Sounds have beginnings, middles, and endings. A sound can have a low-pitched part and a high-pitched part, and this is not just a matter of some object's being low-pitched at one time and high-pitched at another.

This intuitive philosophical picture of sounds as particulars, not properties or qualities, finds empirical support from research on audition. According to our best understanding of the central task of auditory perceiving, sounds are the individuals that ground the grouping and binding of audible qualities.

Perceiving sounds requires discerning coherent and significant streams of auditory information from an intertwined set of signals bound up with irrelevant "noise". Albert Bregman (1990) likens this problem, which he calls *auditory scene analysis*, to determining the number,

size, and location of pebbles thrown into a lake by observing just the motions of a pair of handkerchiefs moved by the waves that travel up two narrow channels dug at the lake's edge. Hearing, as we experience it, is made possible in information-rich environments by the auditory system's ability to sort through the complex information available at the ears and extract cues about significant items the environment contains. The experienced result is a set of distinct, temporally extended sounds heard as generated in the surrounding space. Audition accomplishes this by grouping or bundling audible qualities into distinct auditory perceptual "objects" or "streams". A set of assumptions and grouping principles for auditory perceptual items (auditory objects or streams) enables us to associate correctly the *low pitch* with the *soft volume* and *faraway location*, and at the same time to group correctly the *high pitch* with the *loud volume* and *nearby location*, without mixing things up into a garbled "sound soup" of *high pitch*, *nearness*, *soft volume*, *low pitch*, *loud volume*, and *distance*. Our ability to group correctly the qualities of auditory perceptual objects or streams grounds our ability to discern complex individual sounds in the environment on the basis of information arriving at the ears. Auditory scene analysis amounts to sound perception precisely because the auditory system invokes principles founded upon assumptions that capture genuine regularities in the world of sounds.

The auditory system solves the problem of auditory scene analysis by segregating the auditory scene into separate sound objects or streams characterized by complexes of pitch, timbre, loudness, and location. This answer, in effect, turns on the auditory system's treating the auditory objects or streams in question as particular individuals. First, auditory objects or streams bear pitch, timbre, and loudness and thus serve as the locus for property binding. Second, discrete auditory objects may be represented as distinct both at a single time and across time. That is, distinct sounds can be heard as simultaneous, and successive but qualitatively similar sounds need not be identified. Third, as the term "stream" indicates, they last through time and persist by having duration, and may be represented to persist even through masking noise. Finally, auditory perceptual objects or streams regularly survive changes to their properties though time, as the sound of a spoken word or waning siren demonstrates. These considerations strongly indicate that auditory objects or streams are particulars that ground audible property grouping and binding, auditory attention, and figure-ground distinctions. Awareness of an auditory object or stream constitutes awareness of a sound, an audible particular.

Finally, sounds have sources. Although we commonly experience sounds as sounds of something -- we hear the sound of a car, a bell, or a dog -- that does not imply, in the first instance, that their sources bear or possess the sounds. We might experience a sound without experiencing its source, and sounds might appear to outlast their sources. Sounds, it seems, are produced or generated by their sources. Ordinary objects and happenings cause sounds. Properties and qualities, on the other hand, are not commonly understood as standing in causal relationships to their bearers.

These arguments show that we do not regard sounds merely as repeatables that account for the dimensions of similarity among other items. Rather, sounds are distinct particulars that bear similarity and difference relations to each other based on the complexes of audible qualities -- the properties of pitch, timbre, and loudness -- to which their identities are tied. Sounds have identity, individuation, and persistence conditions that require us to distinguish them from properties or qualities of the objects and happenings that produce sounds.

Identifying sounds with properties has a defect that in my view cannot satisfactorily be repaired. The defect is a failure to account for the essential temporal characteristics of sounds. Property bearers may instantiate and persist through the loss and gain of properties and qualities, while properties, qualities, and their instances exhibit quite different temporal characteristics. This serves as an important indication that sounds are not just properties things gain and lose. The way sounds persist and have duration distinguishes them most sharply from the traditional secondary quality understanding implicit in much philosophical work on sensation and perception. Once appreciated, the temporal characteristics of sounds present the greatest theoretical obstacle to a perceptually tractable and phenomenologically plausible account of sounds along the contours of the property model.

All of this is not to say that no account of properties could make sense of the particularity and temporal character of sounds in a way that dealt with auditory grouping and binding through time. A trope theorist, for example, might capture the particularity of sounds by understanding sounds as particularized complexes of particularized pitch, timbre, and loudness complexes bearing particularized temporal relations to each other. The success of the theory of sounds, however, should not rest on such a controversial theory about the metaphysics of properties. My claim is that given the apparent particularity of sounds, which is required to capture certain aspects of how we perceptually individuate sounds, and given the temporal characteristics of

sounds, including duration and change, the property model assumed by both the traditional secondary quality view of sounds and Pasnau's more recent account is ill suited as a perceptually realistic account of the metaphysics of sounds. Abandoning that model frees us from a host of cumbersome and weakly motivated metaphysical commitments. This points the way to a richer and more nuanced understanding of auditory perception and its objects.

Sounds as waves

The standard philosophical understanding of sounds, of which I have been critical, has not gained broad popularity. The science of acoustics has taught that sounds are waves. We learn early on that sounds are longitudinal pressure waves that travel from a source to our ears and that these waves are the proximal causes of auditory experiences. The sound just is the wave train leading from source to subject.

Just what the customary wave view of sounds amounts to metaphysically is somewhat obscure. One way to characterize the wave is as a pattern of pressures at each point in the surrounding medium over time. This interpretation makes the wave a complex property of the medium that evolves through time. On the version of the secondary quality view that ascribes sounds to the medium, pressure patterns are candidates for the categorical bases of dispositions to produce auditory experiences. This proposal, however, is a version of the property understanding of sounds and faces just the problems that stem from treating sounds as repeatable properties instead of particular individuals. As an account of the metaphysics of sounds it makes little headway. There are, however, other promising ways to develop the view that sounds are waves. If the wave view is plausible as a view about what sounds are, then *the wave* in question is a particular that persists and travels through the medium.

First, waves stand in causal relations. Waves are produced or generated by their sources. Sound waves are the causal byproducts of the activities of objects and interacting bodies and have among their effects the motions of resonating bodies and the auditory experiences of hearers.

Second, the wave bundle responsible for the experience also has spatial boundaries. It is characterized by a wavefront that propagates outward from the source, and its spatial extent depends on when the wave-generating activity ceases and the last pressure disturbance brings up the rear. Even when the waves rebound from a reflecting surface, spatial boundaries may remain

intact, though altered. Furthermore, these spatial boundaries are perceptually significant. For example, the onset of periodic pressure differences at one ear is assumed to share a cause with their onset at the other ear, despite a delay. The spatial boundary responsible for differential onset is critical for auditory localization.

Third, the waves propagate or travel at a speed determined by the density and elasticity of the medium. In 20 degree Celsius air at sea level, we say that the speed of sound waves is 344 meters per second (1497 m/s in water; 6420 m/s in aluminum).

Finally, waves are capable of surviving changes to their shape and to other properties and qualities. A wave's form and amplitude may change as it propagates, resulting in different heard attributes, but the wave persists throughout.

Such spatially bounded, traveling particulars are in certain respects surprisingly object-like. They can be created; they have reasonably defined spatial boundaries, but persist through deformation; they survive changes to their locations and other properties; and they are publicly perceptible. To be sure, they make peculiar sorts of objects: their capacity to overlap and pass through themselves makes them stranger than most everyday objects. Though this may be a mereologically interesting problem, it seems to pose no fundamental obstacle to viewing wave bundles as in some, perhaps minimal, sense object-like.

Another important qualification to this object-like nature is that waves are dependent particulars. Sound waves depend for their existence on a medium. Their survival conditions differ from those of the medium, and they depend on different bits of the medium at different times, but without an elastic medium no sound waves exist. It is likely that lots of other things are dependent particulars, too, like tables and chairs and anything else not identical with its constituting matter. This seems to pose no obstacle to viewing the waves as object-like.

The dependence of waves on a medium is significant for a different reason. In light of the awkward fit of understanding waves as object-like particulars, the dependence points to an alternative take on the wave bundle altogether. The wave is in an important sense something that *happens to* the medium. The wave is not just a parasitic item passing through the medium; it constitutes a dynamic occurrence that takes place within the medium. The existence, propagation, and boundaries of the wave depend on processes that occur within and essentially involve a medium, so to highlight the medium dependence of the wave and its attributes is to highlight the wave's event-like characteristics. It is more plausible to think of *the waves* the wave

conception of sound identifies as the particular sounds not as the object-like bundle, but instead as a variety of event that takes place and evolves in the medium through time.

Whether or not the wave view of sounds can accommodate it, the event-like construal is far more plausible as an account of *sounds* than the object-like construal. Features central to how we conceive of object-like particulars, in contrast to time-taking particulars like happenings and events, make for poor characterizations of sounds. One telling point already played a key role in rejecting the property understanding and delivers a central desideratum in theorizing about sounds. An account of sounds should capture the fact that the qualitative profile of a sound over time is crucial to its being the sound that it is, as we recognize in the difference between the sounds 'protect' and 'protean'. But it is an intuitive feature of the way we perceive and perceptually understand objects that they persist by enduring through time, as opposed to perduring by having numerically distinct temporal parts at different times. That is, we intuitively think of objects, as opposed to time-taking particulars, as being wholly present at each time at which they exist. That is what led Thomson to say of perdurantism, "It seems to me a crazy metaphysic -- obviously false" (1983: 210). And that is why the perdurantist must motivate the view with philosophical considerations. This fact about the way that objects appear to persist does not apply to events and other time-taking particulars, which intuitively have parts that exist and take place at different times. In particular, it does not apply to sounds as we perceptually individuate them, since sounds simply are not candidates for being entirely present at a given moment. Sounds, instead, are things that occur over time.

Now, if objects do perdure, in contrast to the intuitive way we perceive and understand them, then the difference between events and time-taking particulars and objects may be just a matter of degree. If so, sounds are quite a distance from the end of the continuum occupied by tables, chairs, and even persons. In any case, I do not want my account of the metaphysics of sounds to hinge essentially on a discussion of how objects persist. What is clear is that sounds differ in important respects from ordinary objects in their ways of extending through time.

My goal has been to point out that the widely accepted wave view is not completely clear either from a metaphysical standpoint or as a theory of sounds. The understanding of waves as event-like particulars is the most promising way to develop the view that sounds are longitudinal compression waves. That work seems worthwhile because the view that waves are dependent, spatially bounded, event-like particulars that persist and travel from their sources outward

through the surrounding medium captures many of our commonly held beliefs about sounds. But the model of sounds as waves, like the traditional philosophical model of sounds as properties, has important shortcomings that make it unsuitable for a philosophical theory of sounds.

It is a strength of the wave view that it counts sounds as particulars that persist. But a theory of sounds should identify not only the ontological kind to which sounds belong, but also just where in space and time sounds exist. The wave account's problems stem primarily from its implication that such particulars exist or occur in different parts of the medium over time. The claim that sounds travel, however, turns out to be an unnecessary and, indeed, undesirable commitment for a theory of sounds.

The locations of sounds

Any realist account of sounds should say just where in space and time sounds exist. As with property accounts, other theories may differ in where they locate the sounds. If sounds are waves, and waves are events, sounds are located throughout the medium and travel in the sense that their position changes from one time to another. At one time the waves are there but not here; at another time the waves are here but not there.

But hearing, like vision and probably unlike olfaction, is a locational modality. Hearing furnishes information about the locations of objects and events in the surrounding environment. We learn on the basis of hearing not just that a plate has broken, but also something about where to look for the mess. Though hearing lacks the fine spatial resolution of vision, audition presents information about the relative locations of audible events and objects.

Hearing furnishes information about the locations of objects and events in the surrounding environment by presenting sounds themselves as located. Sounds seem to be located not only in a particular direction, but also at some distance. Auditory researchers refer to this phenomenon as *extracranial localization*. One of the most active areas of research into locational hearing seeks to explain the mechanisms that ground the experiential sense that sounds occur at particular locations around us and do not just seem to be located, for example, at the ears.

Why say that sounds themselves seem to be located? First, the data of psychological research supports this claim. In *Spatial Hearing*, Blauert (1997) 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. (38-9)

The spatial information conveyed in audition, however, is not just directional. 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. (45)

Blauert notes that although 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" (45-6). This is representative of the intuitive and widely accepted view among auditory researchers that hearing informs subjects about the locations of sounds in egocentric space.

This view also is apparent when we compare it to alternative phenomenological descriptions. Sounds do not ordinarily seem in auditory experience to travel. Imagine hearing a sound that seemed to be generated across the room, and that subsequently seemed to move toward your head like the auditory analog of a missile. You probably would try to duck out of the way. The experience of such a traveling auditory particular would be quite unlike your ordinary experience of sounds, which seem to be located at a distance in some direction.

We often, however, describe sounds as *coming from* their sources, and not as being at or near their sources. My auditory missile example illustrates that sounds do not auditorily seem to travel toward us from their sources. Sounds also do not seem to be nearby (at the ear) but to have come from somewhere else, as a breeze is felt on the face as having come from the left. Headphone listening illustrates the contrast. Ordinary headphones noticeably lack distance or externalization cues, though they support directional hearing.[2] Sounds therefore do not seem to

[2] Expensive headphones that retain the cues required for externalization exist. Such headphones require custom measurements to determine the effects of the pinnae on incoming sound waves to calculate the individual *head related transfer functions* (HRTFs) that proper externalization requires. See, e.g., Carlile (1996).

come from their sources in any spatial sense of *coming from*. The sense in which it is correct to say that sounds seem to come from their sources must be a *causal* sense. Sounds seem produced or generated by their sources.

The claim that sounds are phenomenologically located in the environment, at a distance in some direction, grounds an important fact about locational hearing. It is clear that that we gain information about the locations of items and happenings around us by means of audition. Furthermore, this locational information is perceptually available to us in audition -- we can act upon and form beliefs about the locations of things in the environment just on the strength of auditory experience. Since sounds are the immediate objects of auditory awareness, awareness of a sound and its audible qualities must furnish or bear locational information about sound sources. But sounds do not seem to come from their sources in a sense that includes travel from those sources, and sounds do not seem to come from their sources in the sense that they seem to be nearby but to have come from the source. Sounds seem to come from their sources only in the sense of being produced or generated by those sources.[3] So, hearing sounds themselves as located makes possible one's audition-based awareness of the arrangement of everyday things and happenings, and it grounds perceptual beliefs about their locations.

I have argued that sounds seem located and that sounds seem to travel only if their sources do. Sounds in this way mediate auditory perceptual access to the locations of things and events in the environment. Unless we are subject to a systematic illusion of spatial location in audition, a theory of sounds must locate the immediate objects of hearing at a distance from perceivers, in the neighborhood of their sources. Not only do we sometimes get the locations of sounds wrong in hearing if sounds are not distally located and relatively stationary, we almost never perceive a sound to occupy its true location. If the phenomenological claim is correct, and if auditory experience is not systematically illusory with respect to the locations its objects, then sounds do not travel through the surrounding medium, and the wave model fails.[4]

---

[3] Nudds (2001) discusses at length the perception of the generation of sounds by sources.

[4] Pasnau (1999) argues that locational hearing is incompatible with a wave account of sounds that does not attribute widespread spatial illusion to audition.

Duration

Locational hearing is not all that is mistaken if sounds propagate through the medium. The illusions multiply. The wave-based understanding of sounds is unable satisfactorily to account for a critical dimension of sounds and auditory experiences. It, too, fails to capture the temporal characteristics of sounds.

Perceiving the durations of sounds is clearly an important part of auditory perception. Sounds inform us about happenings in and states of our environment, and part of what they inform us about is how long those happenings last. I learn through hearing when the coin stops spinning, when the fridge starts up and shuts down, and how long the car idles in the driveway. I experience how long the nine-year-old who lives next door practices violin each afternoon -- I sometimes wish the sessions had shorter durations. If sounds are spatially bounded particulars that travel through the medium, what in fact I experience when I take the sound to have duration, however, is not the duration of a sound at all. Rather, my encounter with a *spatial* boundary of a sound leads to my enjoying an auditory experience while the sound passes. On later encountering the far boundary of the sound, I experience the sound to end. Whether the wave is an object-like particular that passes by, or an event-like particular that unfolds at different places in the medium over time, domino-wise, my experience of the sound is caused by the spatial parts of the sound wave bundle as it passes. I do not experience the lifetime of an object-like entity or the duration of an event other than my own sensing. Apparent duration perception results from encounters with the spatial boundaries of sounds, according to the wave view. This means that each time I hear a sound, I mistake an experience of the spatial boundaries of a sound for an experience of the duration and temporal boundaries of that sound. The experienced duration of a sound is therefore nothing more than a form of crude projective error: I mistake the duration of an experience alone for the duration of the thing I am experiencing. Duration perception, too, is a wholesale illusion if sounds are waves.

Perhaps you are willing to live with the illusion to preserve the common scientific view. So suppose the experience of a sound's duration is an illusion. Since experiencing a sound mediates our awareness of sound-producing events, and, in particular, since experiencing the duration of a sound mediates awareness of the duration of a sound-producing event, your awareness of the durations of sound-producing events is mediated by illusory awareness of the durations of sounds. We have no reason, however, to doubt that awareness of the durations of

sound-producing events is veridical. Such awareness regularly grounds true perceptually-based beliefs. It follows that this case constitutes an instance of veridical mediated awareness that is mediated by an illusion.[5] This complication strikes me as the most important negative consequence of a commitment to illusory sound duration perception, since it is arguably among the primary functions of auditory perception to inform us about the temporal characteristics, including the durations and patterns of change, of happenings in our environment. The account of sounds as waves entails that we do not hear the durations of sounds and that our justification for believing that the violin practice lasted 45 minutes cannot come just from hearing because what we experience is an illusion. These consequences result from the claim that sounds construed as waves travel through the medium.

The other important consequence is that our ways of perceptually individuating and tracking sounds through time are wildly misguided. If sounds persist and travel in the manner of the waves, then our perceptually-based estimates of the lifetimes and survival conditions of sounds all are incorrect because the waves may continue to exist long after the sound has seemed to cease. It is simply a mistake according to the wave account to state that 'Time Is on My Side' by the Rolling Stones is three minutes and two seconds long if the song is the sounds.

I have claimed that the traveling wave view of sounds runs into problems with duration perception, since it makes the perceived duration of a sound a systematic illusion. The problem lies with saying that the sound -- what you most immediately hear -- is the bundle of waves that passes. Suppose we omit the claim that sounds travel. Because it is a central fact about pressure waves that they travel through a medium, we must then abandon the suggestion that sounds are waves. I contend that the illusions of location and duration warrant doing just this.

Sounds, I claim, are located roughly where we hear them to be: at or near their sources. The sound does not travel as do the waves. The waves, however, are causally intermediate between the sounds and the auditory experiences of perceivers. The waves bear or transmit information about sounds through the medium, and thus furnish the materials for auditory

[5] It is important here to keep in mind that the mediatedness in question is of a sort to which the subject has access. It is not, for instance, the kind of mediatedness in question when we say that hearing is mediated by activity in the cochlea or auditory nerve.

experience. Sounds are stationary relative to their sources. If sounds are stationary events, then the auditory experience of location does not involve a systematic and pervasive illusion, and audition-based beliefs about the durations of sounds are for the most part true.

Sounds as events

The wave understanding of sounds gets several things right. According to the best version of the account, individual sounds are particulars that can be counted and quantified over, and possess a range of attributes and qualities. Sounds need not be repeatables or properties ascribed either to ordinary objects or to the medium. It recognizes that sounds are temporally extended occurrents with temporal parts and durations and counts sounds as persisting particulars capable of surviving change. Under its best interpretation, sounds are event-like particulars. A wave-based understanding, however, is unable to capture correctly the temporal characteristics of sounds and the nature of our perceptual acquaintance with sounds that extend through time. In short, it mistakes the lifetime of a train of sound waves in an environment for the duration of a sound.

But the claim that sounds are particular *events* captures important truths about sounds and meets defining desiderata for a theory of sounds and the objects of auditory perception. Sounds, intuitively, are happenings in one's environment. We speak of sounds, like lectures but unlike colors, shapes, and tubas, as occurring, taking place, or lasting. Sounds also stand in causal relations. They are caused by ordinary events like collisions and vibrations, and give rise to reverberant vibration, auditory experiences, and recordings. According to the standard account of causation, causal relata are events. Sounds have straightforward temporal boundaries that circumscribe durations, but, like events and unlike objects, their spatial boundaries are less obvious. Sounds, in addition, appear to tolerate colocation or overlap with other sorts of things and events. A sound might occupy part of the same region as a fiddle or a bowing. Sounds, that is, appear to relate to space and time in ways characteristic to events. Understanding sounds as events of a certain sort amounts to a powerful framework for a satisfactory account of both the metaphysics of sound and the objects of auditory experience.

There is one caveat. The critical features of the theory of sounds should not turn on some one account of the metaphysics of events. I would like the theory of sounds to be reasonably

neutral on the nature of events and viable no matter what events turn out to be.[6] One might even hold it against a theory of events if it fails to capture facts about sounds. So, within reason, whatever events turn out to be, sounds should count as events. I think there is good chance for this, though once we get down to the detailed theory of sounds, some decisions will turn on just what is the right account of events. I want for now to operate with an intuitive conception of events as potentially time-taking individuals -- happenings that may or may not essentially involve change. Events as I wish to understand them are immanent or concrete individuals located in space and time.

Sounds, among the events, are akin to processes or activities. Sounds are not instantaneous events, but require time to unfold. Some sounds -- such as spoken words, birds' calls, or an eighth-note at C-sharp -- may lend themselves to treatment as performances or accomplishments with a certain natural trajectory toward completion.

So, sounds are events located at a distance from their perceivers. They occur at or near their sources, and travel only if their sources travel. Sounds have durations and are capable of surviving changes to their properties and qualities across time. Sounds stand in causal relations to the activities of objects and events that are sound sources, and they fulfill the causal requirement on any account of their veridical perception. Sounds thus occupy distinctive causal roles.

Which distal events are the sounds? Consider the case of a tuning fork struck in air. The striking of the fork makes or causes a sound in virtue of the oscillating arms of the fork disturbing the surrounding air and creating regular compressions and rarefactions. However, since sounds do not travel through the medium, but remain stationary relative to their sources, the sound does not travel as do the waves. Since sound waves that reach the eardrums cause auditory experiences, sounds must be causally intermediate between ordinary, everyday events and traveling sound waves. Since waves bear and transmit information about sounds, sounds cause waves. And since sounds indicate something about the events and happenings that occur in an environment, ordinary objects and happenings cause sounds. Recall that what you perceive as the duration of a sound is in fact the duration of the process of sound wave production. Since the

[6] Candidates include, for instance, theories stemming from Davidson (1970), Kim (1973), Galton (1984), Lewis (1986), and Bennett (1988), among others.

event in which sound waves are produced occupies a role causally intermediate between ordinary collisions or strummings and subsequent sound waves propagating throughout the medium, this event plays a centrally important part in developing the theory of sounds. My claim is that such events are strong candidates for the particular events that are the sounds.

Consider the tuning fork. The sound, I propose, is the event of the tuning fork's disturbing the medium. According to this way of articulating the proposal that sounds are events, particular sounds are events of oscillating or interacting bodies disturbing or setting a surrounding medium into wave motion. This event occupies the appropriate causally intermediate role between the everyday events that cause sounds and the compression waves that travel through the medium bearing the marks of sounds and producing experiences. If a sound just is an object or interacting bodies' disturbing the surrounding medium in a wave-like or periodic manner, then sounds do not travel through the medium, but remain stationary relative to their sources. A sound unfolds over time at a location determined by the sound source. Though it does not travel through the medium, however, it necessarily involves a medium. If sounds are the immediate objects of hearing, such disturbing events are the best candidates for the sounds. Its creating the disturbance constitutes the tuning fork's sounding.

According to this account, sounds are particular events of a certain kind. They are events in which a moving object disturbs a surrounding medium and sets it moving. The strikings and crashings are not the sounds, but are the causes of sounds. The waves in the medium are not the sounds themselves, but are the effects of sounds. Sounds so conceived possess the properties we hear sounds as possessing: pitch, timbre, loudness, duration, and spatial location.

This distal event understanding of sounds counts among its greatest strengths the resources to capture convincingly the conditions under which sounds are identified and individuated. The disturbance event account individuates sounds primarily in terms of their causal sources and their spatio-temporal boundaries. A given sound particular has a unified causal source and must be spatially and temporally continuous throughout its entire history. A change in causal source, or a spatial or temporal discontinuity, suffices for numerically distinct sound particulars. Qualitative resemblance, however, is neither necessary nor sufficient for numerical identity of a sound. A temporally seamless transition from one instrument's playing a C-sharp to another instrument's playing a C-sharp involves numerically distinct sounds of the same sound type, since it involves different disturbance events. Qualitatively similar sounds with

numerically distinct sources are the same sound in nothing stronger than a qualitative sense. Temporally discontinuous soundings from the same source likewise are at most qualitatively identical since they involve different medium disturbance events. But when a single instrument seamlessly shifts from playing C-sharp to playing B, only its state of sounding changes. There is still a single sound event of which each note instance is a part, and so each note instance is part of a single continuous sound. A sound can extend over considerable time and might change a great deal qualitatively. It may at times be loud and low-pitched; at times it may be soft and high-pitched. If the causal source remains the same and the disturbing is spatio-temporally continuous, it may remain a single sound.[7] Difficult cases for the spatial and temporal criteria, such as a tele-transported or time-traveling trumpeter, may of course arise. These cases should be decided by appeal to whether the causal source criterion is satisfied. When the causal source is numerically identical, spatial and temporal continuity from the point of view of the source may obtain and resolve the question in favor of identity.

None of this rules out that there might be complex sounds comprised of distinct sounds from a number of sources arranged either across space, over time, or both, as when an orchestra plays. Complex sounds might even include periods of silence. Consider the sound of a song or of a spoken sentence. Complex sounds, however, are complex events constituted by many distinct sounds or disturbance events, and some principle of unity must exist. There may be many different justifiable ways of counting sounds in these kinds of cases, but ways of counting complex sounds are intelligible because they invoke complex event types or complex sound universals. The ways of counting or individuating sounds may differ depending on one's purpose. Understanding the metaphysics of music or of speech sounds differs from developing a metaphysical account of environmental sounds because the kinds and complexity of the sound events of interest to each enterprise differ. It is striking, however, that disputes over individuation principles for sounds and disagreements about the number of sounds one has heard mirror disputes about individuating or counting events themselves -- disputes that are notoriously difficult to resolve. This makes it a virtue of the event model of sounds that it leaves room for

[7] A significant and sharp qualitative change may suffice for distinct sound particulars in absence of a temporal gap when it diagnoses a different medium disturbance.

disputes concerning how many sounds have occurred, since it inherits that feature from questions and uncertainty about counting and individuating events.

Sound-related phenomena: interference, echoes, and Doppler effects

The discussion so far leaves unresolved a host of questions about pervasive sound-related phenomena. The familiar wave model is fantastically successful at explaining the experiential impact of effects such as constructive and destructive interference; transmission through interfaces and barriers; echoes and reflected sounds; and the Doppler effects. Divorcing the sounds from the waves traveling in a medium means the distal event proposal, or any account that locates sounds at or near their sources, owes equally explanatory accounts of these phenomena and the related contents of perception.

The event model surpasses the wave view's success at convincingly accounting for such phenomena. The distal event account claims that sound waves transmit information about the sounds. It therefore can explain interference, transmission, echoes, and the Doppler effects as wave phenomena that have little to do with sounds themselves. For example, the situation in which destructive interference among the waves from two different sounds creates a "silent spot" at a node where the summed amplitudes of waves cancel is not just a place where there is no sound. Instead, wave interference creates places from which one cannot hear the two sounds that exist in the surrounding space. Because information about sounds is transmitted through a medium by pressure waves, and because waves behave as they do, from such a "silent" node it is as if there are no sounds around. The situation from that place mimics locally the situation, with respect to waves, in which no sounds exist to be heard in one's surroundings.

What about echoes? According to theories on which sounds are waves, an echo is a sound that travels through space and rebounds from reflective surfaces. No such story is available if sounds do not move with their waves. According to the event account I have proposed, the experience of an echo is not a second encounter with a traveling sound at a later stage in its career. An echo experience instead is a second, illusory, experience of the original primary sound. One enjoys a second experience of the original sound event thanks to the way sound waves travel and rebound from reflecting surfaces. The second experience, however, includes illusions of space and time. The echo experience presents the sound as located where it is not (at the reflecting surface), and though the sound heard is past, the echo experience presents it as

occurring *now*. This temporal illusion, however, is no more troublesome than the minimal temporal illusion in ordinary hearing, or for that matter in vision. Sounds heard and objects seen are heard and seen as they *were* due to the delayed arrival of information about them. The delay in hearing is greater than that in vision, as we easily confirm at a fireworks display.

Explaining the Doppler effect is one of the event account's strong suits. According to the wave account there actually are two Doppler effects. When a source travels toward a stationary subject, individual wave peaks compress to yield a higher frequency and higher perceived pitch. Since the frequency of the wave is higher than if the source were stationary, and since pitch is tied to frequency, the pitch of the sound itself is higher than when the same source is stationary. If, however, a subject travels toward a stationary source, the subject encounters more wave peaks per unit of time and falls prey to an illusion of increased pitch. Understanding sounds as events located near their sources, however, yields a unified explanation of source-motion and subject-motion Doppler effects. Both source motion and subject motion produce illusions of altered pitch thanks to how waves transmit information about sounds and excite auditory experiences. In neither case do the qualities of a sound change due to relative motion of source and subject. Rather, a sound merely seems to have altered its pitch thanks to such relative motion. The event view thus captures the way experienced pitch depends upon the motions of subjects and sound sources. As with interference and echoes, Doppler effects are perceptual effects that result from our encounters with waves; none involves the sounds themselves.

Concluding remarks

The foregoing discussion illustrates that the event view furnishes the materials for an explanatorily robust understanding of sounds and their perception. The key is that sound waves transmit information about sounds but are not identical with the sounds. Waves are the proximal stimulus to audition but are not themselves the objects of auditory perceptual experience. This account relies on a model of auditory perception that differs in important respects from the received understanding of hearing as involving awareness of sounds constituted by perceptible patterns of pressure difference throughout a medium. The medium according to the event account is a necessary condition on the perceptibility of a sound, but the medium cannot satisfy the constraints that must be met by the proper objects of auditory perception.

What consequences does this account have for theorizing more generally about perception? Sounds, I have argued, are not among the traditional secondary or sensible qualities because sounds are particular individuals. Pitch, timbre, and loudness, however, provide auditory analogs of color and other sensible qualities. But sounds are not ordinary objects, and sounds are not even object-like particulars. The event view therefore challenges the simple understanding according to which perception reveals just ordinary objects and their attributes. The event view of sounds thus entails that more variety exists among the immediate objects of perception than many modern views acknowledge. If sounds are the immediate objects of auditory awareness, and sounds are events, then audition involves unmediated awareness of events. That is, in the sense discussed earlier, awareness of a sound is awareness of an event that is not mediated by prior awareness of some object and its states or changes. In fact, any auditory awareness of the activities of ordinary objects is mediated by awareness of sounds. Events figure into the immediate contents of audition according to this account.

I have aimed to demonstrate how thinking about perception and the natures of its objects is made richer by attention to audition. Such attention contributes, more broadly, to understanding how to reconcile the manifest and scientific images of the world. The guiding suggestion has been that the tyranny of the visual undermines a complete understanding of perception and the things we perceive. We are likely to miss the most interesting and distinctive features of sounds and audition if we remain bound to the model of vision. Just as we miss what is most striking about vision and its objects if we neglect spatial features, appreciating audition and the nature of its objects requires taking seriously their temporal characteristics. Just as ordinary visual objects are essentially spatial, sounds are essentially temporal. Traditional visuocentric ways of understanding the objects of experience simply fail to capture what is most interesting about sounds.

I have argued that taking the temporal features of sounds seriously shows that sounds are neither secondary qualities as the traditional philosophical outlook has it nor waves as the common scientifically grounded view suggests. Sounds are bearers of audible qualities, and waves, which facilitate audition, are the causal byproducts of sounds. Everyday events like collisions and vibrations cause sounds whose locations remain stationary relative to their

sources. Sounds do not travel. Sounds are particular events whose locations and durations are, when things go well in hearing, as they seem to be.[8]

References

Bennett, J. (1988). *Events and Their Names*. Clarendon Press.

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

Bregman, A. S. (1990). *Auditory Scene Analysis: The Perceptual Organization of Sound*. MIT Press.

Carlile, S. (1996). *Virtual Auditory Space: Generation and Applications*. R. G. Landes.

Davidson, D. (1970). "Events as Particulars." In *Essays on Actions and Events*. Clarendon Press, 1980.

Galton, A. (1984). *The Logic of Aspect*. Clarendon Press.

Kim, J. (1973). "Causation, Nomic Subsumption, and the Concept of Event." *Journal of Philosophy* 70: 217-36.

Lewis, D. (1986). "Events." In *Philosophical Papers, Volume II*. Oxford University Press.

Locke, J. (1975). *An Essay Concerning Human Understanding*. P. Nidditch, ed., Clarendon Press.

Locke, J. (1823). *Elements of Natural Philosophy*. In *The Works of John Locke*, volume III. Printed for Thomas Tegg, London.

[8] O'Callaghan (2007) develops and defends this account of sounds as events.

Nudds, M. (2001). "Experiencing the Production of Sounds." *European Journal of Philosophy* 9: 210-29.

O'Callaghan, C. (2007). *Sounds: A Philosophical Theory*. Oxford University Press.

Pasnau, R. (1999). "What is Sound?" *Philosophical Quarterly* 49: 309-24.

Thomson, J. J. (1983). "Parthood and Identity Across Time." *The Journal of Philosophy* 80: 201-20.