Sounds and Events

Casey O'Callaghan

To appear in
Sounds and Perception: New Philosophical Essays
M. Nudds and C. O'Callaghan
Oxford University Press.

Abstract. I argue that sounds are best conceived not as pressure waves that travel through a medium, nor as physical properties of the objects ordinarily thought to be the sources of sounds, but rather as events of a certain kind. Sounds are particular events in which a surrounding medium is disturbed or set into wavelike motion by the activities of a body or interacting bodies. This Event View of sounds provides for a unified perceptual account of several pervasive sound phenomena, including transmission through barriers, constructive and destructive interference, and echoes.

1 What is a Sound?

Sounds are public objects of auditory perception. When a car starts it makes a sound; when hands clap the result is a sound. Sounds are what we hear during episodes of genuine hearing. Sounds have properties such as pitch, timbre and loudness. But this tells us little about what sort of thing a sound is—which metaphysical category it belongs to. This is the question I wish to answer.

2 Three Theories of Sound

Locke held that sounds are properties of bodies. More specifically, he held that sounds are secondary qualities: sensible qualities possessed by bodies in virtue of the "size, figure, number, and motion" of their parts, but nonetheless distinct from these primary attributes.[1] Robert Pasnau has recently proposed an account according to which sounds are physical properties of ordinary external objects.[2] On what I will call the *Property View* an object "has" or "possesses" a sound when it vibrates at a particular frequency and amplitude. Pasnau follows Locke in claiming that sounds are properties of objects, though he reduces sound to the primary quality that is the categorical base of Locke's power, i.e., that of vibration or motion of a particular sort.

The received view of auditory scientists and physicists is quite different. It holds that a sound is a disturbance that moves through a medium such as air or water as a longitudinal compression wave. Vibrating objects produce sounds, but sounds themselves are waves. When we hear sounds we do not immediately hear bodies or properties of bodies; we hear the pattern of pressure differences that constitutes a wave disturbance in the surrounding medium.

The common interpretation of Aristotle is that he held a very similar view. *De Anima* II.8 says that "sound is a particular movement of air,"[3] which seems to indicate that Aristotle held a version of the received view, or as I will call it, the *Wave View*. We can, however, take our interpretative cues from other passages in the same chapter and arrive at a view that has certain advantages over the other two theories and will be at the core of the alternative I will develop. At 420b13, Aristotle says that "everything which makes a sound does so because something strikes something else in something else again, and this last is air." So, a striking causes or makes a sound when it happens in air. The sound itself is a movement. But the sound need not be the motion *of the air itself*. Instead it may be the event of that medium's being disturbed or moved. The idea is to treat 'movement' as the nominalization of

[1] Locke, *Essay* II.viii.
[2] Pasnau (1999) and (2000).
[3] Aristotle, *De Anima* 420b10.

a transitive verb and focus on constructions like 'x moves y' instead of 'y is moving'. "*For sound is the movement of that which can be moved* in the way in which things rebound from smooth surfaces when someone strikes them,"[4] then means that sound is the air's being disturbed by the motion of an object. A sound is not motion, but the act of one thing moving another. This is not the Wave View that most attribute to Aristotle,[5] but the beginnings of the *Event View* of sound.

According to the Event View, 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 as we shall see, spatial location. When all goes well in ordinary auditory perception, we hear sounds as they are.

3 Locatedness and the Wave View

According to the Wave View, sounds are waves: a particular sound is a train of waves that is generated by a disturbance and moves through the surrounding medium. But this is not how things seem. When we hear a sound, we hear it to be located at some distance in a particular direction. In ordinary cases sounds themselves, not merely their sources, seem to be located distally. Auditory scientists call this phenomenon '*externalization*'.[6] Sounds are not perceived, however, to travel through the air as waves do. They are heard to be roughly where the events that cause them take place: 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."[7] If auditory experience is not systematically illusory with respect to the

[4] Aristotle, *De Anima* 420a20. My italics.

[5] See Pasnau (2000).

[6] Gelfand (1998) refers to this phenomenon as 'extracranial localization': "Sounds heard in a sound field seem to be localized in the environment," p. 374. See also Blauert (1997).

[7] "Feeling That Witnesses Need a Hand, Police Offer One," *The New York Times*, October 16, 2002.

perceived locations of sounds, then sounds are not waves, since they are not perceived to be where the waves are.[8]

The argument depends on a phenomenological claim. Sounds are perceived to have more or less determinate locations. When we hear a clock ticking, the sound seems to be "over there" by the clock; voices are heard to be in the neighborhood of speakers' heads and torsos; when a door slams in another part of the house, we know at least roughly where the accompanying racket takes place. I mean that we experience sounds, in a wide range of cases, to be located at a distance from us in a particular direction. When we do not, as when a sound seems to "fill a room" or "engulf" us, the sound is perceived to be "all around," or at least in a larger portion of the surrounding space. Hearing a sound located "in the head" when listening to earphones is another sort of sound location perception, albeit a touch odd.[9]

Often, however, it is natural to describe sounds as "coming from" their sources: we ask where the buzzing sound is coming from and wonder whether the sound of the cougar came from ahead or behind. 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 auditory spatial perception.

How are we to take talk of sounds' being heard to "come from" a location? Do sounds seem to *come from* locations outside the head, or do they seem to have relatively stable locations outside the head? 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. 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 ordi-

[8]Cf. Pasnau (1999).

[9]Gelfand (1998) refers to this phenomenon as '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," p. 374.

nary experiences of sounds, which present them as stationary relative to the objects and events that are their sources.

Perhaps 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 like listening with earphones: direction without distance. Earphone listening differs from ordinary hearing not just in where sounds seem to come from, but also in where sounds are heard to be. Imagine feeling *where the fan is*. Since sounds seem to come from sources in a sense that includes distance as well as direction, and not travel, the best sense to make of sounds' seeming to come from particular locations is that they have *causal sources* in those locations.

Given the phenomenological facts, the degree to which auditory location perception is illusory or misleading should follow from a theory of sound. No theory should make the fact of location perception a wholesale illusion, though individual instances of location perception might mislead about the actual locations of sounds. Thus, I might correctly hear a stereo speaker's sound as located at the speaker itself; but I might undergo an illusion as of hearing the sound to be located five feet to the right of that speaker. In both cases, I correctly perceive the sound to have a location, but the experience is inaccurate in the second instance. Occasionally, sound location perception is to some degree anomalous, as when sound seems to be all around in a reverberant room, when it seems to be in the head during headphone listening, or when the sound seems to be "behind" a jet plane overhead. Whether and when a sound can literally *fill* a reverberant room, be *inside* the head of a subject, or be *behind* an airplane will depend upon one's theory.

The phenomenon of locatedness spells *prima facie* trouble for the Wave View. Sound waves pervade a medium and move through it at speeds determined by the density and elasticity of the medium. Yet we neither hear sounds as air sloshing around the room nor as moving roughly 340 meters through the air each second. Sounds are perceived to be relatively *stationary* with respect to their sources. The sound of a moving train seems to move only insofar as the train itself moves. When the train stops moving, so does its sound.

The trouble for the Wave View is serious. Since sounds are heard as having stable distal locations, either the sound is not identical with the sound waves, or we misperceive sounds in one important respect. If the sound is identical with the sound waves, the situation is not that we sometimes misperceive sounds, as when a sound ahead is heard to be behind; rather, we *systematically* misperceive the locations of sounds. That is, we hear the locations of *all* sounds incorrectly since we *never* hear a sound to move just as wavefronts do. Since sounds are among the things we hear, we should take the phenomenology of auditory experience seriously when theorizing about what sounds are. If the phenomenon of locatedness is not systematic misperception, then sounds are not sound waves.

The Wave theorist might reply,

> The immediate objects of auditory perception—what we hear—are waves. Sounds just are waves. Waves and their properties are the causes of perceptions of pitch, loudness, and duration; however, we hear these qualities to be located at the place where the waves originate, i.e., at their source. Sounds seem to be where their sources are, and to this extent, auditory perception is illusory. But this illusion is a beneficial one, given our interest in sound sources as constituents of the environment. It is no surprise that we hear sounds to be located where distal objects and events are.

The Wave theorist's response avoids the conclusion that sounds are not identical with waves by accepting that we are subject to wholesale illusion in one salient aspect of auditory experience. The strategy is to assuage concern about the location illusion by providing another candidate for bearer of the spatial properties and by highlighting the illusion's potential benefits. Notice the tactic: by invoking the location of the *source*, the Wave theorist avoids assigning potentially problematic locations to *sounds*.

But an account that locates perceived instances of pitch, timbre, and loudness with their sounds is preferable, all else equal, to one that convicts auditory perception of systematic illusion about the locations of its objects. In part, the case against the Wave View depends on whether there exists an alternative view that captures the locatedness of sounds while matching or surpassing the Wave View's success at

providing a unified explanation of other sound-related phenomena. Part of the task of this essay is to develop such an alternative.

Before going forward, we must first consider: Can we eliminate the location illusion from the Wave theorist's account entirely? A final promising approach again rejects the phenomenological claim as it stands. Instead, it says that we hear sounds to have pitch, timbre, loudness, and duration, 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 properties entirely. We do not mistakenly perceive the locations of sounds, we simply fail to perceive their locations.

This is the Wave theorist's best way to avoid the dilemma. She says sounds are not heard to have locations, they are heard to have located sources. The picture is this: sounds are waves; waves have sources; sounds are heard to be generated by their sources, but not themselves to have locations; only sources are perceived to have locations. This description provides a compelling account of the phenomenology that is consistent with the Wave View. Unfortunately for the Wave theorist, it fails. To see why it fails we need to consider just how audition furnishes perceptual information about the locations where sounds are generated.

Hearing provides information about ordinary objects and events around us—notably, information about where those things are and occur. (Try *not* to turn your head toward a book dropped behind you.) The response we are considering is that we hear objects and events as located by means of the sounds they generate. For the Wave theorist, the basic audible qualities are qualities of sounds, and sounds are waves. Thus, waves have the audible qualities. But we cannot *hear* just non-located audible qualities and located objects, *full stop*. This would amount to a precarious perceptual situation. How could hearing non-located qualities provide perceptual information about sound source locations?

One way is for locational information to be encoded temporally, for example, by time delays between waves reaching the ears. However, since we are auditorily *aware* of the locations of things and happenings—hearing is spatial—this information must be conveyed somehow in conscious perception. At the basic level of awareness,

audition presents just complexes of pitch and timbre with loudness and duration, so an auditory experience that conveys information about the locations of material objects and events must do so by means of one's awareness of these basic attributes. Temporally encoded location information is manifested through one's experience of pitch, timbre, and loudness.

For an experience of 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 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 auditory experience.

A distinction can thus 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 seem to come from their sources only in a causal sense, and since auditory awareness of location must occur by means of awareness of audible qualities, hearing sounds and their qualities as located is required in order to perceive or form judgments about 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. If the Wave View is correct, the location illusion remains.

If the phenomenological claim is an accurate description of the experience of sounds, and if it is true that in order to perceive the locations of sound sources, audible qualities and sounds themselves must be perceived as located, then either the Wave View attributes widespread illusion to auditory perception or the Wave View is false and sounds are not simply waves in a medium. Short of accepting and explaining the illusion, the Wave theorist's best strategy is to impugn my description of the phenomenology. She should say that sources, not sounds themselves, are heard as

located. This requires rejecting the argument that in order to hear sound sources as located, sounds must be heard as located. I know no simple route around the dilemma for the Wave theorist shy about claiming that we fail to perceive, or systematically misperceive, the locations of sounds.

I am convinced that the phenomenological claim is correct as it stands and that sources are heard as located only if sounds are. So I need a way to avoid this difficulty. My theory must not imply that sounds move through the air. The Property View is tailor-made to capture the phenomenology of locatedness. It, however, falls to a separate objection.

4 The Argument from Vacuums

The Property View says that sounds are properties of things like bells, tuning forks, and whistles; more specifically, sounds are the vibrations of material objects. The view entails that sounds are roughly where we perceive them to be. Unfortunately, the Property View also entails that sounds can exist in the absence of a transmitting medium. That is, sounds can exist in a vacuum (just as things can have colors in the dark), since all that is required for an object to have a sound is that it vibrate in the right way.[10] Nevertheless, we have good reasons to believe that the existence of a sound requires a medium. If there can be no sounds in vacuums, the Property View is false.

In Berkeley's first dialogue between Hylas and Philonous, Hylas argues against the Property View in favor of the Wave View by deploying the *Argument from Vacuums*. It begins with the premise that a bell struck in water or air makes a sound, but in a vacuum it does not. Hylas concludes that sound must be in the medium.

> PHILONOUS. Then as to sounds, what must we think of them: are they accidents really inherent in external bodies, or not?
>
> HYLAS. That they inhere not in the sonorous bodies, is plain from hence;

[10] It also appears to entail that there does not exist a *causal* relation between a source and its sound.

> because a bell struck in the exhausted receiver of an air-pump, sends forth no sound. The air therefore must be thought the subject of sound. [The sound which exists without us] is merely a vibrative or undulatory motion in the air.[11]

The argument is:

1. A bell struck in a vacuum makes no sound.
2. So, sound does not exist in the absence of air.
3. So, air is the subject of sound (i.e., the Wave View is true).
4. The Property View is false.

Notice a few things. If a bell struck in a vacuum makes no sound, then sound does not exist in the absence of air (or some other medium) and the Property View is false. But it does not follow that the Wave View is true. Air might be required for the existence of sound without itself being the subject of sound. Even if its first premise is true, the Argument from Vacuums does not establish the truth of the Wave View. Room exists for an alternative theory of sounds according to which no sounds occur in vacuums.

Furthermore, the first premise must be established if it can be used against the Property View. Why say there are no sounds in vacuums? Hylas baldly assumes there are not. We would like to have some reason, preferably independent of an explicit theoretical commitment, for denying (or affirming) that sounds exist in vacuums.

A first pass: "When the bell is struck in a vacuum we know there is no sound because none can be heard by any ordinary creatures; if a medium is added, we can hear it, so the sound must require the medium." Problem: The fact that no sounds are ever heard in the absence of a medium shows only that a medium is required for there to be veridical perception of sounds. It does not show that a medium is necessary for there to be a sound.[12] This is the Property theorist's wedge in the Argument from

[11] Berkeley (1975), pp. 171–2. Also quoted in Pasnau (1999), p. 321.

[12] Berkeley, of course, had reason enough to conclude that there are no sounds in vacuums, since he accepted that nothing exists unperceived.

Vacuums. Suppose one strikes a bell in a vacuum chamber containing a (hypothetical) perceiver. The perceiver can hear nothing. Our problem is that without a theory of sounds we are unable to confirm whether or not there is a sound. Barring a declared theoretical commitment, how do we decide if the bell makes a sound?

Talk of vacuums and sounds might end here, until we have chosen among competing accounts of the metaphysics of sounds. We might, on the other hand, simply bar any view that permits sounds in vacuums, on the grounds that it is too much at odds with common sense. Neither is required; good reasons suggest that sounds cannot exist in vacuums, whether we can confirm it or not.

Perhaps because the bell struck in a vacuum is not a possible object of auditory experience, it does not make a sound? Though I see no good reason to deny that there are sounds beyond the ken of perception, this argument gets us closer to what we are looking for. A sound, if anything, is the bearer of the properties of pitch, timbre, and loudness. Suppose we could establish that there is neither pitch, nor timbre, nor loudness when the bell is struck in the vacuum. We could then reasonably conclude that there is no sound. The bell struck in a vacuum has no sound because it has none of the qualities necessary for the existence of a sound.

The sound of a bell seems to have different qualities when the bell is struck in air and water, and different ones yet in helium and liquid mercury. When the very same striking event occurs in a vacuum, it is inaudible. If a sound exists in a vacuum, it must have some definite pitch, timbre, and loudness. What loudness, for example, does it have? The loudness it would have been heard to have if it had been surrounded by water? The loudness it would have been heard to have if it had been surrounded by air? A decision here will to a significant extent be arbitrary and will not reflect the relevant ways in which the loudness of a sound depends upon the medium in which it is generated.

The Property theorist might hope that ideal or standard conditions for perceiving the "true sounds" of things can be formulated as they can for colors.[13] If so, the pitch, timbre, and loudness of a sound in a vacuum are just those it would appear to have

[13]Pasnau (1999) appeals to just such a hope, p. 322.

in ideal or standard conditions. There is, however, an important disanalogy between colors and sounds. If colors depend on the reflective properties of a surface, then daylight or white light is *normatively significant* in a way that less-than-full-spectrum lighting is not. Once reflected, full-spectrum incident light carries information about how much light a surface reflects or emits at each wavelength across the entire visible spectrum. But the way in which daylight or white light might be counted as ideal for revealing the true colors of things finds no analog in sound. Neither air nor water nor helium does a substantially better job divulging the subtle vibrations of an object in the way that full-spectrum light reveals the reflective properties of a surface. If there is no ideal or normatively significant medium in which to hear the true sound of an object, and if the qualities of a sound depend upon the medium in which it is generated, then it is doubtful whether the object vibrating in a vacuum has a pitch, timbre, or loudness. It is therefore doubtful *whether there is any sound*.

Might we say that sounds are properties that objects have only when in the presence of a medium, and thereby save the Property View? The sound property assigned to the object must in this case depend upon the specific properties of the medium surrounding it in order to avoid the objection raised above. The sound differs when the medium differs. This is no longer the Property View. It is a relational view involving object and medium which is closer to the truth about sound, not the view that sound is an inherent property of objects.

The medium dependence of audible qualities shows that we are justified in drawing a stronger conclusion than that some necessary condition for sound perception is missing in a vacuum. A necessary condition for there *to be a sound* is missing. If sounds do not occur in vacuums, the Property View is false.

Moreover, we showed earlier that the Wave View entails systematic illusion about where sounds are. The Event View, however, is a natural alternative that attributes the right locations to sounds and does not entail that sounds exist in vacuums. The Event View is that particular sounds are events in which a medium is disturbed or set into wave-like motion by the movement of a body or interacting bodies. These *disturbance events* take place where we perceive sounds to be, and, because no medium

is present to be affected, a vacuum contains no sounds.

5 The Event View

Particular sounds are events.[14] Sounds take time and involve change—at a minimum they begin, and usually they end. A number of qualitatively different stages or a single tone of uniform loudness may compose a sound. The sounds are the events in which a medium is disturbed or changed or set into motion in a wave-like way by the motions of bodies. Events such as collisions and vibrations of objects cause the sound events. Among the effects of sounds may be sound waves propagating through a medium and the auditory experiences of perceivers. Medium-disturbing events are what we hear to have particular pitch, timbre, loudness, and location. A body counts as in a state of sounding—making a noise—just in case it is in the midst of generating or causing a particular sound. Whenever there is a sound there is a sounding.

The tuning fork struck in air is a simple case. The striking is an event that "makes" a sound in virtue of the process by which the arms of the fork oscillate and create regular compressions and rarefactions in the surrounding air. Its creating the disturbance constitutes the tuning fork's sounding. The event of the tuning fork's disturbing the medium is the sound. We perceive this sound event to have a constant pitch and timbre, a duration, a location, and diminishing loudness. In contrast, the sound of an owl's call is a more complex event characterized by a temporally extended pattern of changing pitch, timbre, and volume. Each call sounded is an event that consists in the owl's lungs and syrinx disturbing the surrounding air in a given pattern. The tuning fork and the owl alike are recognizable by the sounds they create.

Auditory perception also makes us aware of events in our environment. We learn by audition how the furniture is arranged and when it is being moved. How is this possible if sounds themselves are the events that we hear? The Event View says that a

[14]My theory of sounds as events should be relatively insensitive to what particular theory of events is the correct one (e.g., Lewis (1986), Davidson (1970), Kim (1973)). Within reason, whatever events turn out to be, sounds should be events. Accordingly, I wish to work with the intuitive notion of events as particulars which take time and may or may not essentially involve change.

sound is an event whose cause is the event heard to "have" or "make" the sound, and implies that the sound and its cause are in close spatio-temporal proximity, since we might treat the location of the disturbing event at a time as the surface of interaction between the object and the medium.[15] When we hear the sound of a glass breaking, that sound is an audible event constituted by the fracturing glass affecting the air. The breaking of the glass causes the medium-affecting event that is the sound event. The medium-affecting event is near the breaking event, but the two do not occur in just the same space-time region. That the two sorts of events occur close to each other, however, does not sufficiently explain why we are aware of *sound generating* events in auditory perception. A sound also carries qualitative information that can be used to identify its generating event after perceivers learn to associate the sound with the cause. The sound's pattern of pitch, timbre, loudness, and duration indicates that a glass has broken; the location of the sound points us in the direction of the mess.

So, there is the event of an object or substance setting a medium into periodic motion. This is a sound. The kind of motion depends on the form and makeup of the object or substance, what it does to disturb the medium, and the physical characteristics of the medium itself. The sound event has a location and a pattern of pitch, timbre, loudness, and duration. There are also the generating events that cause sounds and the objects that are said to make a sound in virtue of instances of their sounding.

Sounds are individuated along three primary dimensions: causal source, spatio-temporal continuity, and qualitative change. Intuition is sometimes silent, but we do have implicit in our practices principles for saying when sounds are the same or different. The Event View captures these principles. To count as the numerically

[15] Cf. Bennett (1988):

> "The 'location' of an event is its spatiotemporal location, i.e. where and when it occurs. ... A zone may be sizeless along one or more of its dimensions: ... some fill spatial volumes and presumably others occupy only planes, lines, even points," p. 12.

same sound particular, a candidate must have the very same token causal source and be spatially and temporally continuous throughout its entire history. If either the causal source changes or there is a spatial or temporal discontinuity, we say that there are different relevant sound particulars—a temporally seamless transition from a trumpet playing B-flat to another trumpet playing the same note counts as involving two different sound tokens. The sense in which the sound from a single trumpet is different when it seamlessly goes from playing a B-flat to an A is that the trumpet's state of sounding is different. Perhaps different sound events correspond to these different states of sounding at the two times. Still, there is one sound event of which each note instance is a part, and in this sense both are parts of a single continuous sound. Such a sound might extend over considerable time and space and change greatly in its qualitative characteristics. At times it may be loud and high-pitched and at others it may be faint and low, but as long as it has the same causal source in terms of its generating event or object, and is spatio-temporally continuous, it counts as the very same sound particular.

It follows that numerically distinct instances of sounds that fall under the same qualitative characterization are not *the very same sound* in any sense stronger than qualitative identity. Temporally discrete sounds from the same causal source and spatially separated sounds generated by different sources can at best be different instances of the same qualitative sound type. Philosophers sometimes do, however, speak of performances of songs and symphonies as events that are tokens of sound types, despite the fact that they need be neither temporally nor spatially continuous—they may incorporate periods of silence and multiple sources. We can say the same of bird calls. But these are complex events that involve patterns of individual sound events when they occur. Distinct sound particulars are arranged to comprise a whole that may require or allow for discontinuities of various kinds. The ontology of music and complex sound universals enjoys its own vast literature. What I want to point out is that the Event View is capable of capturing the ways in which we take sounds to be individuated. The principles I have mentioned may be disputed; but there is often obscurity about how events are to be individuated. The Event View, in that case,

predicts—correctly—that there is a certain amount of obscurity and arbitrariness in our verdicts concerning how many sounds we have heard.

The Event View is a natural way to avoid the objections posed to the Wave View and Property View. Particular soundings have audible locations determined by where the medium-disturbing process occurs. Sounds, then, move through space in just those ways we expect them to, for example, when a train passes in the distance. The subject-directed missile-like sound does not ordinarily arise. The Event View also accounts for what we learn about sound from the Argument from Vacuums: we are justified in claiming that a medium is necessary for there to be a sound. Since a medium is required for there to be a medium disturbance, there is no sound in a vacuum. The Event theorist maintains that sounds are neither entirely in the surrounding medium nor simply properties of objects. If the arguments against the prevailing views are compelling, the Event View is a theoretically cogent solution.

Several lines of objection force elaboration of the Event View. The Event View provides for natural accounts of several phenomena that pose difficulties for any theory of sound.

6 Transmission

So far, I have said that a sound is an event of a medium's being disturbed or set into motion in a particular way by the activities of an object, body or mass. But this seems too lenient, to allow too many sounds. Consider the following two forms of objection.

> **First Form**
>
> Suppose you are underwater and hear the sound of something that happens in the air above, say, the striking of a bell. The Event View seems to imply that there is a sound at the interface of the air and water since indeed there is a medium-affecting event there. The air, a mass or body, sets the water, a medium, into motion. This is phenomenologically inaccurate. We do not hear the sound to be at the surface of the water; we

hear it to be above in the air.

Second Form

The preacher outside is loud. When I shut the window I do not hear him as well. The window muffles the sound. Nevertheless, the window also sets the medium inside the room into motion. According to the Event View, is the sound located at the windowpane? We do not hear it as being there—the sound still seems outside.

In both forms, sound waves generated in one medium pass into another kind of medium. The first describes travel across a single interface; the second involves travel through a solid barrier. In each, at the relevant interface—the air-water interface in the first case and the window-room interface in the second—the motion of a body disturbs the medium it adjoins. Yet since we do not ordinarily take ourselves to hear sounds at such places, intuition has it that no sound occurs at either the interface or the barrier. Must the Event theorist count these events as sounds?

The problem of transmission is not unique to the Event View. Each of the views canvassed faces a version of the objection. The Property View is in roughly the same straits as the Event View. The Property View implies that the sound is a property of the air mass in the first case and the windowpane in the second, since each vibrates at a particular frequency and amplitude. Even the Wave View, on which sounds are waves, faces a dilemma. What is the *source* of the sound? Is it the bell or the air-water interface? The preacher or the window? Each is in a sense the cause of the waves "in the medium" in which the sound is heard. An acceptable version of the Wave View must acknowledge that we perceive locations in auditory perception, even if these are the locations of sound *sources*. Just as the Event theorist needs to say which events are the sounds, the Wave theorist must say which things count as sources of sounds. Though the problem is not unique to the Event View, the Event theorist owes an account of sounds and transmission.

The Event theorist's options are: (a) deny there is a sound where transmission occurs and explain why the Event View does not entail that there is; (b) accept that sounds accompany transmission events and reconcile this with the intuitive de-

scription of the experience. Contrary to first appearances, option (b) is somewhat attractive. Ultimately, however, this response with its burgeoning world of sounds is unsatisfactory. It strains the imagination to suppose that a multiplying of sounds occurs each time sound waves travel across an interface or through a barrier. Our accounting should be more sober.

Suppose we deny that a sound occurs when a "new" medium is disturbed by a pre-existing sound wave. Option (a) suggests a attractive way to conceive of the perceptual situation. We say that the interface or barrier distorts our perception of the primary sound's location and qualities, not that we perceive a secondary sound with its own location and set of qualities that is caused by the primary sound. A single sound exists above the water or outside the window, but one may not have an ideal experience of that sound if impediments to perception intervene.

This picture is more accurate from a phenomenological standpoint. We have a perceptual bias toward the locations of sound generating events of the everyday sort such as doors shutting and ocean waves breaking. We hear the sound created by the striking of the bell above water and the sound of the preacher proselytizing outside the window. Events of transmission occur when the waves from one sound event cause motion in an object or body that is passed on to another medium. We do not hear events of transmission or indeed anything at their locations when we hear a sound beyond an interface or through a barrier.

The language of this distinction points to a theoretical solution compatible with the Event View. To speak of a sound or of a sound wave as *generated* by a source implies that the sound or the wave is caused by and distinct from the event that brought it about. The idiom suggests that neither the sound nor the wave exists prior to an event of generation. In contrast, the idiom of *transmission* suggests the passing along of a wave disturbance that already exists. Indeed, the physics of sound wave generation differs from that of sound wave transmission. During generation, something which is not itself a sound wave produces a sound wave; during transmission, sound waves *travel through* an interface or barrier as sound waves. Sound events involve the active production of pressure waves, transmission events do not.

When a transmission event causes a medium disturbance of the sort that seemed to pose trouble for the Event theorist, that event depends for its existence upon a prior sound event. The distinction between events in which sound waves are introduced into an environment and those in which sound waves are transmitted is natural and based on the events' roles in a regular causal network. Being a sound is a matter partly of occupying a particular causal role. A central feature of the causal role distinguished by how we speak, one supported by the physics, is that sounds are events caused by generating events such as collisions, but are not caused simply by waves passing through boundaries and barriers. The medium-disturbing events that are the sounds are the events in which a wave disturbance is *introduced* into the environment by the activities of some material object, body, or mass. Medium-disturbing events in which a prior sound's waves are passed on or transmitted into a different medium are not in any ordinary sense sounds.

Suppose sound waves reach a barrier and induce vibrations in that object. The barrier might then itself generate a sound in addition to the sound whose waves induced the barrier's vibrations. This is not an ordinary case of sound wave transmission, however, and should be subsumed instead under *resonance*. Resonating is sounding since the resonating object actively disturbs the medium, and does not merely passively transmit existing sound waves.

This account appeases intuition. The problem of deciding which of multiple sounds we listen to when sound waves pass through an interface or barrier does not get off the ground. But the innocent picture according to which being a sound is entirely a matter of what happens near the surfaces of objects whose activities affect a medium is threatened. We must adopt a broader perspective that acknowledges the causal relations of several distinct kinds of events. This is not cause for alarm, nor is it a surprise given the organization of sound-related experience. Sounds furnish us with awareness of sound generating events, which are of paramount interest for what they tell us about the world. They tell us such things as how the furniture is arranged and when it is being moved. Transmission events, however, enjoy little utility beyond what we learn through their effects on how we perceive the primary sounds they occlude:

when we perceive a sound as muffled, we learn that a barrier may intervene. Given our interest in ordinary events that take place among material bodies, along with how these events are related to sounds, it is no wonder that the primary disturbances should be distinguished by audible qualities.

7 Destructive and Constructive Interference

As commonly demonstrated in physics classrooms, sound waves *interfere* with each other. Suppose you are in an anechoic room in which two tuning forks tuned to E above middle C are simultaneously struck. As you move around the room, there are places from which you hear the sound to be soft and places from which you hear the sound to be loud; there are places from which you hear neither sound.

This phenomenon occurs because at any time the total pressure at a point in the room equals the algebraic sum of the pressures of all the sound waves at that point. It is therefore possible, when sound waves are out of phase with each other, for the total pressure at some point or in some area to remain constant while separate sound waves pass through that point or area simultaneously. A listener positioned at such a point hears nothing. When sound waves cancel, the interference is *destructive*. When the waves are completely in phase at a point, the total pressure varies with the sum of the components' amplitudes. The sound seems twice as loud as either tuning fork at these points thanks to *constructive* interference. Altering the phase or vibration characteristics of one of the tuning forks may result in *beating*, a periodic variation in perceived volume from a particular point.

Here is the problem. Take the example of complete destructive interference described above. The Wave theorist can explain that you hear no sound from where you stand because the pressure is constant at that point and hence there is no sound. Of course, there are still in a sense two sets of waves passing through that area, though their summed amplitude is zero. So, in a sense, there are two sounds at that point even though none is heard. The Wave theorist does not escape entirely. If, however, by 'the wave' we mean something that depends only on the total pressure at a point,

there is no wave and no sound at the point in question. By contrast, the Event View implies that each tuning fork makes a sound even though you hear neither one from the point of interest. If sounds are not sound waves and the Event View is correct, then you hear no sound at all when there are two. Is the gap a fault line in the Event View?

Interference phenomena do not undermine the Event View. The interference arguments do show that waves carry information about sounds. The Event theorist should not deny this when he says that the sound is not identical with the waves. Waves can be involved in the process by means of which a sound is heard without the sound's just being the waves. The Event View provides an intuitive and compelling alternative to the standard account of destructive interference. The Event View says there are two sounds, two events of a disturbance being introduced into a medium. These disturbances travel as compression waves and may reach a perceiver, where they cause perceptions of the original sound event. Waves obey the principles of interference, and if no variations in pressure exist, no sounds are heard. Ordinarily, a lack of pressure variations indicates the absence of sounds and sound sources. Complete destructive interference resembles the absence of sounds because factors conspire to create nodes where the pressure does not vary. These factors include the spatial arrangement of the two sources, the frequency and amplitude at which the sources oscillate, and the temporal relations among the activities of the sources, i.e., the phase difference of the sources. A perceiver located at a node will hear neither sound, and may believe that no sounds occur. This does not entail that the room contains no sounds. The observer is simply unable to perceive the sounds because of her particular point of view.

That there are indeed two sounds can be confirmed in several ways. One can move to a point where one or the other sound is audible, move one or both of the sources so that the nodes are shifted, alter the phase difference in the vibrations of the two objects to remove nodes completely, or simply remove one of the sources to eliminate interference entirely. These exercises show that each tuning fork makes a sound that can be heard independently of the other in the right circumstances.

Sometimes, however, another sound's presence can interfere with perceiving a given sound. Experience need not reveal from a particular vantage point all the surrounding environment's sounds. What we perceive from a very limited vantage point need not be the entire story about what sounds are around.

The case of constructive interference is very similar. Due to the spatial and temporal relations among events of sounding, a perceiver in the right location may experience multiple sources to have greater loudness than any single source present. This is again the result of the additive properties of sound waves. It is less surprising that the subject's loudness experience should increase in the presence of two sources than that it should decrease, as in destructive interference. Beating is perhaps less intuitively comprehensible, but is also an explicable result of how the source events are arranged in time and space, and of the subject's vantage point on these events.

8 Echoes

The phenomenon of an echo is familiar. You are at a fireworks display in an open field with a single brick building behind you. A colorful bomb's recognizable boom follows on the heels of its visual burst, but a moment later the boom's echo sounds at the brick wall behind the field. This phenomenon poses two potential problems for the Event View. First, is the echo a distinct sound event that occurs at the reflecting surface, or not? Though the echo seems distinct, the brick wall reflects sound waves and does not introduce a disturbance into the surrounding medium, so the Event View appears to have no sound to identify as the echo. Second, does the existence of echoes show that sounds themselves travel and can be re-encountered, and hence, that sounds are not the events I have suggested?

If echoes show that sounds are not events, then the Event View is false. So, the first question pre-supposes a negative answer to the second. I shall argue that echoes do not pose a problem for the Event View, and that once we have secured the correct conception of hearing an echo, the Event View has precisely the right kind of disturbance event on offer: the primary disturbance.

Matthew Nudds has recently argued in the following way that sounds are not events.[16]

1. Newton measured the speed of a sound by measuring the time it took for the sound to travel down a colonnade and back.

2. Hearing an echo is re-encountering a particular sound.

3. Events, unlike objects, cannot be re-encountered.

4. Therefore, sounds are object-like particulars and not events.

If this reasoning is cogent, an echo is a particular sound at a later stage of its continuous career, after it has been reflected.

Nudds also considers and rejects a two-part response: (a) Newton measured the speed of sound waves, but not of a sound; (b) an echo is a *distinct* sound whose qualities resemble the primary sound. Though I reject (b) for reasons I will soon discuss, I do accept (a). Claim (a) is a strong replacement for (1), from which (2) does not follow. But (2) does not follow even from the weaker (1*).

1*. Newton measured the speed of *sound waves* by measuring the time it took for him to hear an echo after hearing the primary sound.

Hearing an echo may not be re-encountering the same sound at a later stage of its career, even if we owe the episodes of hearing to the same sound waves. If (2) ought to be rejected on independent grounds, (1) is false. Since the conclusion that sounds are not events does not follow from the argument reconstructed with the uncontroversial (1*), the Event theorist can then provide an alternative account of echoes. What reason do we have to reject the claim that hearing an echo after its primary sound is re-encountering the same persisting sound particular?

Sounds are essentially extended in time—each sound has a beginning, a middle, and an end. Sounds are not wholly present at each moment at which they exist. Having a qualitative profile *over* time is central to the identity of a particular sound

[16]Nudds (2001). The argument appears on pp. 221–222, and in endnote 15, p. 227.

such as a spoken word or an owl's hoot. One kind of re-encounter we have with sound particulars includes those that occur at later stages "during the completion" of the sound. This morning I heard the loud, high-pitched beginning of the local emergency siren's wail. I then descended into the silent basement for two minutes, after which I emerged to hear the nearly completed sound's fading low-pitched moan. This was a genuine re-encounter with the same particular sound at different times. I experienced a different part of the sound upon each hearing. Now, if Nudds is right, I can also hear a sound *in its entirety* during two (or more) distinct intervals during which it exists, and thereby re-encounter it. That would make sounds particulars that are **extended in time** and that **fail to be wholly present at any moment** at which they exist *and* that **can be experienced in their entirety at different stages of their continuous careers**. But a single particular cannot *continuously* exist throughout an interval of time during which it must begin and end—*entirely*—multiple times. Either sound duration perception—perceiving that sounds begin and end—is illusory, or the claim that sounds are persisting particulars that can be re-encountered in their entirety is false.

However, we perceive the durations of events that produce sounds, such as fingernails scraping across a blackboard, by perceiving the durations of sounds. We do not, in the first instance, perceive the durations of our experiences of sounds. But that is just what we must do if sound duration perception is an illusion that occurs in virtue of encounters with the spatial boundaries of passing sounds. How else could we perceive the durations of events such as blackboard scratchings if sounds did not have the durations we think they have? The trouble comes from thinking that a bout of echo perception is a re-encounter with the same sound particular later in its continuous career.

That claim finds little support in perceptual experience. Hearing an echo is unlike re-encountering a person you have met before; it is unlike glimpsing someone carrying home the vase you saw earlier in a store window, even though the echo has an equally rich qualitative signature. Echo perception does not bear the marks of object-recognition and identification that characterize the experience of material ob-

jects and continuants with relatively stable qualities. An echo seems to be distinct from its primary sound in a way that an object perceived at different times does not. Perhaps this is because particular sounds are often perceived to begin and end, or because one could not imagine continuously perceiving the entire sound as it traveled from source to wall and back. Whatever the explanation of this disanalogy, perceptual evidence fails to bolster the claim that sounds travel and can be heard again as echoes. In the face of the arguments, (2) should be rejected. If echo experiences are not re-encounters of the sort we have with objects, the conclusion that sounds are not events does not follow.

This brings us back to the first problem mentioned in this section. What event is the echo? Is an echo a distinct disturbance event that occurs at the reflecting surface? There are four reasons that together suggest it is not. First, awareness of an echo normally furnishes awareness of the event that made the sound. Hearing the echo of the firework's boom is hearing the *explosion itself* again. A sound is always the sound of something happening. An echo experience, as well as that of a primary sound, discloses those happenings. Second, we do not attribute dispositions to produce sounds with particular audible qualities to the reflecting surface. Third, and importantly, what occurs at the reflecting surface is not the introduction of a disturbance into the surrounding medium. An elastic collision between the surface and the medium occurs, causing the direction of wave propagation to change. Absent is the Event View's characteristic event: the original disturbance of a medium by the activity of a body. The reflecting body need not do anything but redirect pre-existing waves. Finally, an analogy with mirrors is compelling. Mirrors facilitate our seeing the very objects and events that occur in front of them, albeit with distortion of place. Likewise, reflecting surfaces allow us to hear the very sounds that occur in front of them, albeit with distortion of place *and time*, which results from the speed of sound waves. If the mirror analogy is correct, just as there are not distinct visible objects located at the surfaces of mirrors, echoes are not distinct sounds that occur at surfaces that reflect sound waves. Together, these four claims suggest that echoes are not distinct disturbance events that occur at reflecting surfaces.

The picture gestured at by analogy with mirrors is that hearing an echo is hearing the primary sound event over again. This is a re-encounter of a different sort from that rejected above. The sound event occurs only once, say, between t_1 and t_2, and is perceived again between t_3 and t_4 (both later than t_2) because the waves it creates return. The sound neither travels nor returns to the perceiver; the perceiver experiences the same distal event over again because of the way the event's traces travel. Hearing an event that is past is thus like seeing an event that is past. Compare seeing a supernova from across the galaxy. If we could put mirrors in far outer space, we could see the same earthly event twice: once when it happened and once after its traces were reflected.

Why does the apparent distinctness of echoes from primary sounds, which I invoked against Nudds's argument, not tell equally against the claim that when the echo phenomenon occurs, we hear the very same event twice? We hear the event (qua echo) with distortion of location, but our experience as of the echo also occurs *later*. If an echo were an object experienced at a time later in its career, we would expect ordinary object-recognition to occur, given the echo's qualitative similarity to the sound initially heard. In fact, with objects we *count on* this sort of recognition to ground the perceived continuity of our material world. *Capgras Syndrome* is one form of *delusional misidentification syndrome* in which patients suddenly begin to believe that people and objects familiar to them have been replaced by exact qualitative duplicates; this failure of perceived continuity is notable and debilitating.[17] Events and time-taking particulars, however, are tied to a specific time and place when they occur. Though the 2002 World Cup Final might have been located at various times and places, it in fact occurred June 30, 2002, at International Stadium Yokohama, Japan. That very event cannot occur again or elsewhere. Similar events experienced at different times and places are taken to be distinct events. So, if we happen to perceive the *very same* particular event over again, it *should seem* like a distinct event. Since echo phenomenology arises when we hear the very same sound event over again at a later time and different place, precisely what we should expect is the apparent

[17] See Breen, et al. (2000).

distinctness of echo from primary sound.

On this model, we *can* perceive the same sound event twice because of how waves propagate. The situation is something like this. Suppose you hear the sound of the firework. You then travel faster than the sound waves, overtake them, and halt. You now hear the sound again—it seems to be in the same place it was before. We need not say the sound travels, only that the sound waves travel. Because of how information about sounds is transmitted through a medium, you are lucky enough to experience the same sound event over again. The medium disturbance you hear when you hear the sound for the second time is the very same disturbance event you heard earlier. Echo perception is similar. A reflecting surface, however, saves you the trouble of supersonic travel. You pay the price with distortion of location. The Event View nicely captures the correct way to conceive of echoes and echo perception.

9 Concluding Remarks

The Event View replaces the picture according to which sounds fill the air and travel as waves. Instead, sounds are events that occur where objects and bodies interact with the surrounding medium. Sounds are events that take place near their sources, not in the intervening space. Sound waves travel through the air carrying information about these distal events, and are the proximal causes of sound experiences in subjects; sound waves, however, are not sounds. The revision more accurately captures how we experience sounds to be.

The Event View is a natural account of what sounds are that avoids the dilemma concerning where sounds are located. It implies that sounds are distally located and stationary relative to their sources without making them solely the properties of material things. We should not accept the view that sounds are properties of objects because sounds and sources stand in causal relations, and because we have good reason independent of the received view to think that sounds cannot exist in vacuums. The event that the Event theorist identifies as the sound cannot occur in the absence of a medium.

Taking sounds to be particular events of objects disturbing a surrounding medium furnishes a unified picture of what counts as a sound in cases that pose problems for any such theory. Sounds do not occur at barriers where transmission takes place. The phenomena accompanying constructive and destructive interference arise because of the spatial and temporal relations among sound sources and because information about sounds is transmitted by waves. Hearing an echo is hearing with distortions of place and time. The Event View entails no mysteries about sounds and sound experience.

References

Aristotle. *De Anima*. In J.L. Ackrill (ed.), *A New Aristotle Reader*. Princeton: Princeton University Press, 1987.

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

Berkeley, G. *Three Dialogues Between Hylas and Philonous*. In M.R. Ayers (ed.), *Philosophical Works*. London: Dent, 1975.

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

Breen, N., D. Caine, M. Coltheart, J. Hendy, C. Roberts. (2000). "Toward an Understanding of Delusions of Misidentification: Four Case Studies." *Mind and Language* 15: 74–110.

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

Gelfand, S.A. (1998). *Hearing: An Introduction to Psychological and Physiological Acoustics*. 3rd ed. New York: Marcel Dekker.

Kim, J. (1973). "Causation, Nomic Subsumption, and the Concept of Event." *Journal of Philosophy* 70: 217–236.

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

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

Nudds, M. (2001). "Experiencing the Production of Sounds." *European Journal of Philosophy* 9: 210–229.

Pasnau, R. (1999). "What is Sound?" *Philosophical Quarterly* 49: 309–324.

——. (2000). "Sensible Qualities: The Case of Sound." *Journal of the History of Philosophy* 38: 27–40.