1

Introduction: The Philosophy of Sounds and Auditory Perception

CASEY O'CALLAGHAN AND MATTHEW NUDDS

1. Sounds and Perception

'Humans are visual creatures', it is common to observe. Our reliance upon vision is apparent in the way we navigate and react to our surroundings. We fumble in the dark and instinctively turn to look at the sources of sounds. Visual information also occupies a privileged epistemic role, and our language frequently reflects a tight coupling of seeing with knowing. We evaluate *views*, have *insights*, and *see* what is at issue. Perhaps most telling is the greater fear many admit at the prospect of losing sight over any other sense.

Not surprisingly, philosophers investigating the nature of perception and perceptual experience have considered vision nearly exclusively. Philosophical discussions of sensible and secondary qualities have focused upon color and color experience, while debates about perceptual content primarily concern the content of visual experiences.

Until remarkably recently, something similar was true of empirical researchers who aimed to unearth the processes, mechanisms, and principles that explain how we become acquainted with our environments. Driven by the goal of computer vision, vision scientists were among the first to shed sensory psychology's early preoccupation with psychophysics and the measurement of sensations. Empirical work on perceiving and attending to visual objects has since advanced to the point that Brian Scholl (2001: 2) has described it as 'a type of "case study" in cognitive science'. Vision is better understood than any other sense modality.

But humans are not solely visual creatures. Exclusive attention to vision distorts the degree to which we rely on each of the senses to cope with

information-rich surroundings. Recently, interest has grown rapidly in understanding the other sense modalities and sensible features that figure in our capacity to negotiate and understand our environments. Spurred in part by a growing body of rich empirical research, philosophers increasingly have turned attention to tactile, proprioceptive, and kinaesthetic perception; smell and olfactory experience; and aspects of the philosophy of taste (see, e.g., O'Shaughnessy 1989; Martin 1992; Scott 2001; Gallagher 2005; Lycan 2000; Batty 2007; B. Smith 2007). The 'other' sense modalities present challenging new puzzles for the empirical and philosophical study of perception.

No topic in extra-visual philosophy of perception has generated as much attention in recent years as that of sounds and audition. While Strawson (1959) set an early example in *Individuals* by exploring the conceptual consequences of a purely auditory experience, and Evans (1980) responded with a revealing discussion of the requirements on objective experience, the past decade has seen a flurry of work on the nature of sounds and the content of auditory experience. Current research on the perception of speech sounds and spoken language, the experience of music, auditory-visual cross-modal illusions, and the nature of 'auditory objects' promises to impact and advance the philosophy of perception.

More important, however, it signals a departure from the tradition of relying upon vision as the representative paradigm for theorizing about perception, its objects, and its content. While the implicit assumption has been that accounts of visual perception and visual experience generalize to the other senses, nothing guarantees that what is true of seeing holds of touching, tasting, or hearing. Intuitions about critical issues or particular cases might differ in the context of different modalities. While it might seem obvious in the case of vision that perceptual experience is *transparent*, or that space is required for *objectivity*, gustatory and olfactory experiences might tell otherwise (see, e.g., Lycan 2000; A. D. Smith 2002).

Furthermore, resolving certain issues might require examining modalities other than vision. For instance, the debate whether the phenomenological characteristics of experiences are a subset of their representational properties turns on whether visual and non-visual experiences that share representational properties share phenomenological character. Resolving this question depends upon whether it is plausible that all non-visual experiences have representational content, whether visual and non-visual experiences can share representational content, and how best to characterize the phenomenology of non-visual experiences. Given the present state of debate, whether intrinsic properties of experiences constitutively contribute to their phenomenology might only be apparent upon considering

experiences in other modalities and phenomenological differences among modalities.

Even if one's sole concern is vision, examining the other modalities enriches one's understanding of what it is to perceive visually and of how we ought to characterize the phenomenology and content of visual experience. Debates about vision and visual experience are informed by attention to other sense modalities.

Some cases even indicate that one cannot give a complete account of perceiving in any single modality without appreciating phenomena that involve other modalities and without addressing the relationships among the senses. For instance, given an important class of inter-modal effects and cross-modal recalibrations and illusions, the content of vision might in certain respects depend either upon the content of experiences that take place in other modalities or upon amodal content that cannot be characterized exhaustively in purely visual terms. In either case, information associated with another modality impacts experience in vision and helps to determine its content. Whether the relationship between extra-visual information and visual experience is constitutive, merely causal, or entirely accidental, a complete accounting that explains these visual processes and experiences requires understanding of the other senses and the relationships among modalities.

In addition to helping advance familiar debates in the philosophy of perception, the case of sounds and audition reveals new puzzles. One example is whether and, if so, how we hear anything but sounds. For instance, when a door slams, I hear its sound. But I also seem to hear the slamming of the door. The slamming is what motivates me to react. So, while I hear the sound of a door slamming, is it also fair to say that I hear the door itself? If so, how do things other than sounds enter into the contents of auditory experiences, and what is it to auditorily represent a door? Alternatively, are the sources of sounds perceived only indirectly thanks to one's awareness of sounds?

Another example involves the nature of sounds themselves. Traditionally, sounds have been grouped with the colors, tastes, and smells among secondary or sensible qualities. Recently, however, a number of philosophers have argued that sounds are not qualities or properties at all, but instead are events. On this account, sounds are more analogous to visual objects than visible features, in that sounds are the bearers of audible features. This raises a number of questions. If sounds are events, what is it to experience an event in a way that does not depend upon experiencing its participants? Do we experience a source to generate or cause a sound?

The philosophy of sounds and audition also opens new fronts in the philosophy of perception. Considering sounds and hearing forces philosophers

to confront the cases of music and spoken language. Listening to music and perceiving speech provide fascinating examples of hearing's richness and complexity. The possibility of an emotionally engaging temporal art of sounds and the existence of a fluid and flexible communicative medium comprising sounds illustrate the extent to which audition is a significant and central perceptual domain that should not be ignored by the philosophy of mind and perception.

This collection comprises original essays that address the central questions and issues that define the emerging philosophy of sounds and auditory perception. This work focuses upon two sets of interrelated concerns.

The first is a constellation of debates concerning the ontology of sounds. What kinds of things are sounds, and what properties do sounds have? For instance, are sounds secondary qualities, physical properties, waves, or some type of event?

The second is a set of questions about the contents of auditory experiences. How are sounds experienced to be? What sorts of things and properties are experienced in auditory perception? For example, in what sense is auditory experience spatial; do we hear sources in addition to sounds; what is distinctive about musical listening; and what do we hear when we hear speech?

This introductory chapter has three aims. It presents a survey to provide context for the issues discussed in the chapters that follow. It summarizes the main debates and arguments at stake in this volume. And it suggests promising areas for further work, including unsettled questions and topics that remain unaddressed.

2. The Ontology of Sounds

A theory of sounds should identify the ontological kinds to which sounds belong, and it should say what sorts of properties sounds possess. Debates about the nature of sounds have focused upon such questions as whether sounds are mind-dependent or mind-independent, whether they are individuals or properties, and whether they are object-like or event-like. Also, there has been considerable debate about just where sounds are located.

2.1 What Kind of Thing is a Sound?

Sounds are among the things we hear. Auditory experience is directed upon sounds. Sounds, therefore, are *intentional objects* of audition (see Crane 2009). Since it is plausible that sounds are perceived only through the sense of hearing,

sounds commonly are counted as *proper* sensibles of audition. Furthermore, it is plausible to say that whenever you hear something, and whatever you hear, you hear a sound. It is doubtful you could hear something without hearing a sound. Arguably, this is because whatever you hear—such as a collision or a trumpet—you hear it by or in virtue of hearing its sound. Sorensen (Chapter 7), however, disagrees. He argues that we hear silence, which does not involve hearing a sound. Traditionally, nevertheless, sounds are counted among the *immediate* objects of audition.

Given their status as immediate and proper objects of audition, it is not surprising that the nature of sounds has been tied to our experience of sounds. Since at least the early modern era, the predominant view has been that sounds are secondary or sensory qualities. Locke, for one, grouped the sounds with the colors, tastes, and smells as dispositions whose characterization tied them essentially to the experiences of subjects. In the 20th century, some theorists held that sounds are subjective and private and that they mediate auditory perceptual access to the world (e.g., Maclachlan 1989).

Sounds, however, need not be counted as private and subjective given their status as immediate objects of audition if we reject that perception enlists subjectively accessible intermediaries, as do contemporary representationalists along with direct realists and disjunctivists (see, e.g., Tye 2000; Noë 2004; Martin 2006). Sounds then might be experientially or subjectively immediate, which allows either that perception involves no mediators (including representations) at all, or that it requires no experientially accessible but subjective or private mediators.

Sounds might still be grouped with other *perceptible qualities* or *properties*, such as colors, smells, and tastes. For instance, Pasnau (1999) argues that sounds are properties that either are identical with or supervene upon vibrations of things such as bells. On this account, sounds are properties attributed to things commonly taken to be the sources of sounds.

Some recent philosophers have argued that sounds are not properties or qualities, but instead are *individuals* or *particulars*. Rather than *qualifying* or *being properties attributed to* things, sounds are individuals that bear sensible features such as pitch, timbre, and loudness. Sounds on this view are not mere dimensions of similarity.

O'Callaghan (Chapter 2; see also 2007), for instance, claims that property theories do not capture the individuation and identity conditions for sounds. O'Callaghan claims that sounds persist through time and survive changes in ways that sensible qualities and features do not. This raises the question whether sounds are object-like individuals or event-like individuals. O'Callaghan argues that sounds do not simply persist, but have *durations* and commonly are

individuated in terms of the features they exhibit *over time*. For example, the sound of a police siren comprises a certain pattern of changes in audible features over time. The sound of the spoken word 'siren' differs from that of 'silent' in that the two involve different patterns of change through time. So, many sounds are individuated in terms of patterns of features over time. This, and the difficulty of imagining an instantaneous sound, suggests sounds are essentially temporal.

Impressed by the temporal natures of sounds, several philosophers have argued that sounds are *events* of a certain kind. Casati and Dokic (Chapter 5; see also 1994, 2005) identify sounds not with the property of vibrating, but with the *event* of an object's vibrating. O'Callaghan identifies sounds with a closely related but different event. O'Callaghan argues that the presence of a medium is a necessary condition not just upon the perceptibility but upon the existence of a sound, and proposes that sounds are events in which vibrating objects or interacting bodies actively disturb a surrounding medium. This account differs from Casati and Dokic's in three ways. First, sounds are not identical with vibrations. Either they are causal byproducts of vibrations, or they are vibrations only under certain conditions. Second, sounds may result from events such as collisions or strikings in which multiple objects interact. Finally, sounds require a medium and thus cannot exist in a vacuum.

Scruton (Chapter 3; see also 1997) offers a very different kind of event theory of sounds. Scruton rejects the physicalism of Casati and Dokic and O'Callaghan, and argues that sounds are what he calls *secondary objects* and *pure events*. First, on analogy with secondary qualities, sounds, like rainbows and smells, are *secondary* objects of perception. Secondary objects, unlike secondary qualities, are independent particulars or individuals rather than properties or qualities. But, like secondary qualities, they are not identifiable with any physical features or objects. The features of such individuals include just their ways of appearing. Secondary objects are objective, though simple and irreducible. Scruton also claims sounds are *pure* events that do not happen *to* anything and that cannot be reduced to changes to other reidentifiable particulars. Sounds thus lack a constitutive ontological connection with the vibrations or activities of objects we ordinarily count as sound sources. Appreciating the independence of sounds from sources, according to Scruton, is critical to understanding distinctively musical experiences: hearing music *requires* the ability to experience sounds as independent from their physical causes (see Section 4.2 below).

Perhaps surprisingly, none of these accounts constitutively ties sounds to longitudinal pressure waves that pass through an elastic medium such as air or water or metal. Such waves propagate from their sources outward towards observers, have frequency and amplitude, and cause auditory experiences.

According to common sense tutored by science, sounds just are traveling waves.

Several authors in this collection, including Nudds (Chapter 4), O'Shaughnessy (Chapter 6), Sorensen (Chapter 7), and Smith (Chapter 9), endorse theories inspired by the common scientific account. Sorensen, for instance, says, 'Since I identify sound with acoustic waves, I think silence is the absence of acoustic waves' (p. 140). Nudds argues that even though sounds are not identical with waves, they are dependent upon waves. More carefully, he argues that sounds are *instantiated* by waves. According to Nudds, sounds, such as those of words or symphonies, can be instantiated on different occasions and by different waves and frequency patterns. Nonetheless, we may perceptually identify a sound as the very same sound whenever it is instantiated. Nudds thus claims that sounds should be understood either as particularized types or as abstract particulars that are instantiated by the waves. The virtue of this account is that sounds themselves are repeatables, but they are not features of waves, a medium, or objects. This view preserves the intuition that we can make or hear the same sound on multiple occasions while rejecting the claim that sounds simply qualify their sources.

2.2 The Locations of Sounds

One main disagreement between the wave-based accounts of sound such as those of Nudds, Sorensen, and O'Shaughnessy (see also Hamilton, Chapter 8) and source-based accounts such as those of Pasnau, Casati and Dokic, and O'Callaghan (see also Matthen forthcoming) concerns the locations of sounds. The former locate sounds in the medium and imply that sounds propagate and thus occupy different locations over time, or travel. The latter hold that sounds are located at or near their sources and do not travel through the medium—sounds travel only if their sources do.

Debate surrounding this issue draws attention to a substantive constraint on theorizing about sounds and their natures. How we experience sounds to be serves as a prima-facie basis for any account of sounds. This is because, in the first instance, our access to sounds is through auditory experience, and our conceptions of sounds are grounded in experience. An account of sounds should be an account of things it is plausible to identify with sounds as we experience them to be. How our experiences of sounds present them to be thus constrains what account it is plausible to give of the nature of sounds. One way to formulate this experiential constraint on theorizing about sounds appeals to veridicality. An account of sounds should entail that auditory experiences of sounds are for the most part veridical; all else equal, it should not imply that experiences of sounds involve wholesale illusions. So, we might hold

that for any feature sounds are experienced to have, it at least is possible for experience to be veridical in that respect. A weaker version holds that, even if the experience of a sound could not be veridical in all respects, sounds should have at least most of the features we experience them to have. This means that, all else equal, for some feature we experience sounds to have, we should prefer an account that does not ascribe illusion with respect to that feature. We can put the constraint as a slogan: avoid attributing unnecessary illusions.

Distal sound theorists commonly argue that sounds seem in auditory experience to be located at or near their sources. Sounds, they claim, do not seem travel from the source towards your ears, do not under ordinary conditions seem to pervade the medium (perhaps they do under special circumstances, such as in a loud nightclub), and do not seem to be nearby or at the ears. Instead, they claim that sounds auditorily seem to be where the things and events that generate them are located. If we do experience sounds to be distally located, and if sounds are roughly where they seem to be, then sounds do not travel through the medium as wave accounts imply. Distal theorists charge that unless we systematically misperceive the locations of sounds, sounds do not travel through the medium as do pressure waves (Pasnau 1999; O'Callaghan, Chapter 2). In that case, the veridicality constraint means that we should favor the distal view. Hamilton disagrees, and argues that we hear only where the traveling sounds have come from, rather than where they are. A related response is that we hear, veridically, only a subset of the locations of sounds.

The distal theories support an account according to which auditory perception is in important respects analogous to vision. In particular, sounds located at a distance are perceived thanks to a medium (pressure waves) that bears information about them. Sound waves on this account are like the light that conveys information about distal objects and stimulates vision. The physical waves are not the sounds, and the sounds do not travel with the waves, but the waves mediate between sounds and hearers.

On the other hand, some authors maintain that auditory perception differs in this respect from vision. Suppose that in audition we experience a sound that is proximal when we experience it, and that, in virtue of experiencing the sound, we perceive something that is distal. On this account, the sounds heard are located near their perceivers, but they provide information about distal things and events beyond the world of sounds. Such a proximal theory of perceived sounds preserves the metaphysical dependence of sounds upon the sound waves that stimulate hearing. In effect, it locates the sounds we hear (at the time we hear them) at a different stage in the causal chain that leads from source to subject. That causal chain begins with the activities of

things in the environment, leads to wave-like motion in a medium, continues with stimulation of the auditory sense organs, and culminates in auditory experiences. Distal theories locate the sounds we hear at an earlier stage in the causal sequence than do proximal theories.

Since proximal theorists do not wish to say that auditory experiences involve a systematic spatial illusion, they must reject the distal theorists' phenomenological claim that sounds seem in audition to be located at a distance in some direction. Proximal theorists and distal theorists therefore disagree about how best to describe the spatial aspects of auditory experience. Resolving the issue therefore requires a closer examination of spatial audition.

3. Spatial Audition

It would be difficult to deny that hearing conveys spatial information. On the basis of audition, you learn that the barking dog is behind you or that the door to your left has closed. But we can explain in different ways how you learn this. Distal theorists argue that you hear sounds to be located at some distance in a given direction and thereby come to learn about, and perhaps even to hear, the locations of their sources. At the other extreme, subjects might merely *infer* or *work out* information about space and locations from entirely aspatial auditory experiences (O'Shaughnessy, Chapter 6; see also 1957 and Malpas 1965). Smith says: 'Sounds, in general, are hard to place in the spatial world and auditory perception gives us no clues as to where they might occur' (p. 202). The disagreement concerns whether or not audition itself involves perceptual awareness of spatial characteristics, and to what it attributes those spatial characteristics.

Skepticism about spatial audition has been widespread at least since Strawson's (1959) famous claim that a purely auditory experience—in contrast to a purely visual or purely tactile-kinaesthetic experience—would be entirely non-spatial. Strawson claims that a world of sounds would be a no-space world because sounds are not intrinsically spatial. According to Strawson, spatial concepts have no intrinsically auditory significance, and audition's spatial capabilities depend upon its inheriting spatial content from other modalities.

While Strawson's arguments are subject to different interpretations and have been challenged (see Nudds 2001; O'Callaghan forthcoming; Casati and Dokic), they suggest an alternative way to understand how audition grounds spatial beliefs. First of all, not all contemporary proximal theorists wish to deny the vast body of research showing that for perceptually normal subjects with

vision, touch, etc., under ordinary circumstances with multimodal stimulation, hearing itself is spatial. Under such conditions, auditory experience might have spatial content or represent spatial features (see Blauert 1997; Nudds, Chapter 4; Casati and Dokic, Chapter 5), whether this depends upon other modalities or not. Nonetheless, one might claim that we do not experience *sounds* as having spatial features. Nudds, for instance, argues that sound *sources*, rather than sounds themselves, are auditorily experienced as distally located. This accommodates the empirical evidence about auditory localization without accepting that sounds themselves are experienced to be located. On his account, information embodied in sound waves about the locations of sound sources is used to determine and auditorily represent the locations of sound sources without representing sounds as distally located. Such an account might go on to claim that sounds seem to be located at or near the ears, that they seem nearby but to have *come from* some direction, or that they seem to lack spatial features entirely. Therefore, although audition has spatial content, it need not attribute spatial properties, such as distal location, to sounds. Sounds might seem nearby or nowhere, while sound sources seem located at a distance.

Distal theories maintain that information about the locations of sound sources is provided by the audible locations of sounds *at* their sources. In contrast, some proximal theories that attribute spatial content to auditory experiences hold that audition attributes spatial properties to sound sources. Both proximal and distal accounts thus may hold that auditory experiences have spatial content, or that spatial properties are represented in audition. But they may disagree about that to which spatial properties are attributed.

Two things are worth noting. First, the kind of proximal account just described owes an explanation for how audition could represent sound sources as having spatial characteristics without representing sounds as located or as having spatial features. How could sound sources auditorily seem located distally if sounds do not?

Second, in considering where sounds are located, we need to consider where sounds are experienced to be located. This, in turn, leads to a discussion of spatial audition. The facts about spatial audition, including the auditory experience of spatial features of an environment, however, appear to be compatible with the view that we hear sound sources, rather than sounds, as located. Evaluating this alternative to the claim that sounds are distally located thus forces us to consider what audible attributes ordinary objects and events that generate sounds, such as bells and collisions, possess. The proposed account requires that, in addition to the sounds, we are capable of auditorily perceiving the sources of sounds. While distal theories may allow for the

auditory perception of sound sources, their account of spatial audition does not obligate them to do so.

So, there are different options if auditory spatial beliefs about the environment are grounded in spatial audition. First, we auditorily experience distally located sounds, perhaps along with their sources. Second, we hear sounds locally or aspatially, but thereby experience distally located sound sources. Alternatively, one could deny (implausibly, in our view) that auditory experience itself has spatial content.

It is doubtful whether introspection of auditory experience alone could decide among these possibilities (see, especially, Schwitzgebel 2008 for doubts about phenomenological introspection; see also Remez and Trout, Chapter 11, discussed below in Section 4.3). Audition provides lots of useful information about things and happenings that generate sounds. Indeed, one way to individuate sounds appeals to their causal sources. Experiences of sounds thus are closely associated with perceptual information about their environmental sources. Reflecting upon the phenomenology of spatial experience alone may not be decisive without some independent way to determine where we experience sounds to be and whether we auditorily experience sound sources.

The dispute over the locations of sounds thus turns on a family of questions about the content of auditory experience. In addition to sounds, do we hear their sources? Which properties—in particular, which spatial properties—does audition attribute to each? Progress on these issues requires a more detailed study of the content of auditory experience.

4. The Content of Auditory Experience

4.1 Sounds and Sources

Accounts of the content of auditory experience can be sorted into three classes. First, austere views hold that we immediately hear only sounds and their attributes, such as pitch, timbre, loudness, duration, and location. Second, more permissive accounts hold that we might hear both sounds and their sources. According to such accounts, we might hear the sound and hear the bell or its striking. In that case, we also might auditorily experience sounds in some sense to *belong* to their sources. For instance, sounds might be heard as properties or as parts of their sources. Alternatively, sounds might be heard to be distinct from their sources, in which case we also might hear the relations

between sounds and sources. Third, an account could maintain that we hear even things beyond sounds and their sources, such as how things stand in the environment. For instance, in hearing the sound of footsteps I might also hear the enclosed space in which they are being taken.

Deciding among these options poses a methodological challenge. We might appeal to what we say we can hear, or to what we can learn on the basis of hearing. Typically, we say we hear the bird singing as well as the sound it makes. We report learning about the locations of sound sources such as cars or collisions on the basis of hearing. But, with vision, we can say we see that the mail carrier has come on the basis of seeing the pile of mail without being committed to claiming that visual experience represents that the mail carrier has come. So, perhaps we can say that we hear that the bird is singing on the basis of hearing the sound, without saying that auditory experience represents the bird. In general, we need to distinguish what is part of the content of experience from what we learn or judge based upon experience. Though we learn about the sources of sounds on the basis of hearing, appealing to what we can normally come to know on that basis is not an infallible guide to the content of auditory experience.

While we might appeal to the phenomenology of auditory experience to support one or another account of its content, we turned to considering the content of auditory experience in part to avoid relying entirely upon phenomenological introspection. Nonetheless, there are considerations that support thinking that awareness as of sources is an important part of the content of auditory experience. It would be difficult otherwise to explain why we so persistently form beliefs about the sources of sounds on the basis of audition without inference or further assumptions, and it would be difficult to account for the fact that we act on the basis of auditory experience as if we heard sound sources. Reflexively turning to look for the source of a sound or ducking when you hear something coming from behind would make little sense unless you were aware of sound sources. Furthermore, we could make a strong case that your auditory experience as of the sound of a bell would not be veridical if you opened your eyes to see a loudspeaker or a duck.

We might appeal to a general metaphysical view about the nature of perceptual experience, such as a sense datum view (which perhaps favors an austere account) to decide the issue. However, the goals of theorizing about audition and sounds include testing such accounts and learning if they generalize. Furthermore, most contemporary accounts of perception, such as direct realism or intentionalism, are compatible with each of the options.

Another alternative is to appeal to the function of auditory perception and to the kinds of psychological explanations into which auditory content enters.

Though the bulk of laboratory work on audition has used artificial tones in artificial situations, a growing body of work on ecological psychoacoustics appears to provide support for the claim that how auditory processes deal with acoustic information depends in important ways upon natural constraints that amount to assumptions concerning the physical world and properties of sound sources (Neuhoff 2004). For instance, features of *sources*, such as material and size, which determine how they vibrate and disturb the medium, explain dimensions and degrees of auditory similarity and difference that acoustic characteristics alone cannot (see, e.g., Handel 1993; McAdams and Bigand 1993; Bregman 1990). For instance, explaining timbre perception probably requires appeal to features of sound sources (see Handel 1995). This supports a compelling conception of the role of audition as furnishing awareness of the things and happenings in our environments that make sounds.

Perhaps unsurprisingly, then, a prominent theme throughout this collection is that awareness of sound sources is an important aspect of auditory experience. Several contributors here reject the austere claim that we immediately hear only sounds and so must infer or judge what produced them (e.g., Nudds, Chapter 4; Hamilton, Chapter 8; Smith, Chapter 9).

Those who endorse that sources are part of the content of audition do not *just* hold that in addition to hearing sounds, we hear the things that in fact are their sources. Rather, they generally hold some view about the relationship we hear sounds to bear to their sources. Co-location is one such relation (as are other spatial relations). Another possibility is that sounds are heard to be properties of or to qualify their sources. This option is unavailable to those who reject property views of sounds for reasons such as those outlined above. Another possibility is that sounds are heard to be mereological parts of their sources (O'Callaghan 2008). On such a view, sounds might be heard to be parts of events that involve ordinary objects such as bells and whistles. For instance, two cars are involved in a collision, and part of that event is a sound. Hearing a collision by hearing its sound might be akin to seeing a cube by seeing its facing surface (cf. Scruton). A final possibility, according to which sounds are heard to be caused or produced by their sources, perhaps fits best with ordinary thinking about sounds. This requires that we are able to perceive causal relations. It also requires experiencing sources as independent from their sounds, so it remains to explain how we are perceptually aware of sound sources as such.

If audition does involve awareness of sound sources, then audition differs in an important respect from vision. One's auditory awareness of sound sources intuitively is not as *direct* as when one sees those same sources. Thus, even accounts on which we hear distally located sounds, if they also allow that we

hear sound sources, might imply that audition involves a form of awareness of sources that is less direct than visual awareness of objects. This invites a new discussion of the ways in which perception may be direct or indirect that extends beyond the visual case.

It is noteworthy that so many have found it compelling that auditory awareness does not stop with sounds. This contrasts with vision, where fewer have been inclined to say that we see what is causally upstream from the objects, colors, and shapes we visually experience. In the visual case, accounts of 'seeing in' and 'seeing as' and 'metaphorical seeing' commonly are invoked. In contrast, hearing a bell or a bird that makes a sound requires no similar act of imagination.

Considering whether sounds or sound sources are auditorily experienced as located led us to consider whether audition involves awareness of sound sources. While it is not obvious that we auditorily experience sound sources, there are some reasons for thinking that we do. Obstacles remain. What relationship are sounds experienced as bearing to their sources? What features can sound sources be auditorily experienced as having? Why acknowledge indirectness in audition if not in vision? This debate cannot yet settle the question about the locations of sounds. However, it does impact how we should characterize auditory experience, and it raises more questions than it answers about auditory content.

4.2 Music

An account of human auditory perception should accommodate music. Since speech raises special questions that we will address in turn, consider *pure* or non-vocal music. The possibility of an art of non-vocal sounds raises special questions about the nature of musical listening. Does hearing music require a distinctive act of listening? What is aesthetically significant about listening to pure music? This depends upon what is aesthetically significant about music. Since, presumably, we are capable of hearing at least some aesthetically relevant features of music, it also depends on the content of our auditory experience of music. Because it is prima-facie plausible to think that the aesthetic significance of pure music depends only upon sounds in abstraction from the environments or circumstances of their production, however, the case of music contrasts with the case of ordinary audition. This contrast may illuminate non-musical auditory experience. Reflecting on musical listening may, therefore, provide evidence to help resolve the questions about auditory content addressed above.

Is listening to music just a variety of ordinary auditory experience, or is it special? For instance, does musical listening require unique or special capacities or skills? On one hand, music involves sounds and sound sequences,

arrangements, or structures. So, ordinary auditory capacities are needed for hearing music. If music is nothing more than sounds, such capacities should suffice. However, it *is* plausible that one could perceive auditorily without hearing music as such. Animals, for instance, might hear sounds without experiencing music. Musical experience might involve more than just hearing sequences of sounds. But the difference could just be a matter of how one *responds* to one's auditory experience. For instance, music often incites emotions, imaginations, or associations that are triggered by hearing patterns of sounds. Nevertheless, such responses are responses either to a distinctive variety of auditory experience or to particular aspects of one's auditory experience. What are the features of auditory perceptual experience when listening to music that make possible the distinctive experience of music?

Roger Scruton (Chapter 3; see also 1997) argues that musical listening requires hearing in a way that abstracts from one's interest in the environmental sources of sounds. According to Scruton's *acousmatic thesis*, humans' capacity to appreciate music depends upon the unique ability to auditorily experience sounds as detached entirely from their physical causes, or as divorced from the worldly sources of their production. The aesthetic characteristics of music, according to Scruton, are independent from such facts as that individual sounds are produced by an oboe, or a particularly rare oboe, or that a passage requires a high level of skill to perform, or that a performance is live rather than recorded. What matters are the sounds. (Recall, for Scruton, sounds are secondary objects and pure events that are independent of their sources.) This account of musical listening requires that in some sense it is possible to have auditory experiences whose contents include sounds but not sound sources. On a strong reading, listening to appreciate the aesthetic character of music requires auditorily experiencing sounds, without experiencing their sources. That would seem to require an austere, sound-only account of auditory content. One alternative is to deny that musical listening requires austere auditory content and to hold, instead, that musical listening is a matter of *attending* to that which is (independently motivated to be) aesthetically relevant, to wit, the sounds. This modification requires only the capacity to experience and attend to sounds *as* independent from their sources, rather than the capacity to experience sounds without experiencing their sources.

Andy Hamilton (Chapter 8; see also 2007) resists Scruton's acousmatic thesis that musical experience involves awareness of sounds that divorces them from their sources, and argues that attending to sounds as part of the world in which they are produced is an aesthetically relevant aspect of musical experience. Hamilton offers a twofold account on which acousmatic and non-acousmatic listening both provide valuable musical experiences. Hamilton suggests that

features that outstrip sounds, such as the skill of a performer, or the fact that sounds are produced by a performance rather than by a recording, can be aesthetically relevant. Since Hamilton holds that many such features, in addition to sounds themselves, can figure in auditory experience, he argues that auditorily experiencing music involves non-acousmatic experiences. Hamilton thus holds that there is a sense in which we can hear the production of sounds through hearing alone. Sources therefore must enter the contents of auditory experience on this view of musical experience.

But Hamilton also holds that the experience of music is not purely auditory. First, there are aesthetically relevant features of music that we experience through senses other than hearing—including sounds! 'We feel as well as hear sounds' (p. 166), and we see as well as hear the virtuosity of a performance. Such extra-auditory experiences must be non-acousmatic. Moreover, Hamilton doubts whether even acousmatic experience must be purely auditory and thus unimodal. Given multimodal influences that shape perception, listening to sounds in a way that abstracts entirely from their sources, and from other senses, may prove impossible. In that case, multimodal or amodal aspects of perceptual experience may unavoidably infect auditory experience. In that case, even 'purely' auditory experiences of sounds might have non-acousmatic features.

Scruton would simply resist that non-acousmatic aspects of auditory experience are relevant, and he might reject that the other senses matter to our appreciation of music. But, *if* auditory experiences of sounds *unavoidably* have non-acousmatic features, then the acousmatic thesis as stated requires revision. Scruton might comfortably speak of the *aspects* of auditory experience, or of the *features* of sounds, that are aesthetically relevant. This, however, is compatible with rejecting that a special mode of musical listening exists. If musical listening is a unique variety of auditory experience, perhaps it involves a distinctive way of aligning auditory attention. In that case, the skillful act of musical listening could be like an abstract or formal, non-representational mode of looking at paintings or pictures, a way of looking that involves an appreciation of the arrangements of colors and pigments rather than of what is depicted. Arguing that attending to formal features of sounds is the only aesthetically significant way of listening to music raises questions similar to those raised by the corresponding claim about looking.

4.3 Speech

Just as humans, perhaps uniquely, are in a position to hear sounds as music, we also may be unique in hearing sounds as speech. Speech, like music, raises questions about the contents of auditory perceptual experiences. In particular, to what extent do the experiences of hearing speech and of hearing ordinary

environmental sounds share auditory perceptual content? However, the case of speech also introduces complexities that force us to reconsider whether sounds are among the objects of speech perception. Moreover, some researchers even maintain that speech perception is a unique perceptual modality. Thus, the philosophical issues about speech perception concern different versions of the question: Is speech special?

Hearing and appreciatively listening to music involves focusing on acoustical properties of sounds. Perceiving spoken language, however, requires not just hearing sounds, but also grasping that they are sounds of speech. Speech sounds interest us because they bear meaning and communicate linguistic information.

On a traditional account of perceiving spoken language, we hear certain sounds and then grasp their meanings. We auditorily perceive sounds, but we *understand* their meanings. On this account, hearing speech sounds is just like hearing non-speech sounds, except in its effects upon the understanding. Speech sounds cause us to grasp meanings with which they are contingently associated.

Nevertheless, hearing speech in a language you understand differs from hearing speech you do not understand. The difference is not just that in one case, but not in the other, you associate meanings with the sounds you hear. The difference is unlike that between seeing written words you understand and seeing those you do not. The visual experience of the shapes and spacing of letters and words does not change dramatically when you understand them. However, the sounds themselves differ in auditory appearance once you learn a spoken language. You hear pauses, word boundaries, and subtle distinctions in vowel and consonant sounds that you previously did not hear. Understanding a spoken language makes a distinctive difference to the phenomenology of hearing speech sounds.

John McDowell has claimed that understanding a language makes possible the experience of sounds as publicly *meaningful* (1998a, 1998b; see discussion in Smith, Chapter 9). Hearing meaningfulness implies a difference in auditory experience between listening to speech in a language you know and listening to speech in one you do not know. While it offers a richer account of the content of auditory experience in the case of speech perception, hearing meanings does not explain why we experience sounds to have different acoustic qualities once we hear them as meaningful (since meanings lack acoustic characteristics). It also invites us to ask the challenging question: What is the *auditory* experience of meaning like?

Barry Smith (Chapter 9) advocates a more conservative response than McDowell to the traditional account of the roles of audition and the

understanding in hearing spoken language. Smith suggests that, while we understand but do not hear meanings, we do hear more than just the sounds of speech. Smith argues that we are auditorily aware of the *voices* of individual speakers, in addition to the apparent sounds of speech. Awareness of voices, rather than hearing meanings, accounts for our sense of communicative contact with verbal language users. Two features of Smith's account are noteworthy. First, voices play a role similar to sound sources, considered above, in the content of audition. Both are among the things we auditorily experience thanks to hearing the sounds they produce. Smith's account of speech perception thus involves something like hearing sound sources. Second, though it avoids any mystery about the auditory experience of meanings, hearing as of a voice does not by itself account for changes in the experience of sounds and their attributes in a language we understand. Hearing voices can explain the difference that accrues thanks to hearing sounds as speech, but it cannot explain the further difference due to understanding that speech.

So, we might claim that the *contents* of speech perception experiences differ from those of hearing sounds in ordinary non-linguistic audition. Auditory contents when hearing speech might include, as we have seen, meanings or voices. On the other hand, we could attribute the phenomenological difference after learning a language simply to ascribing different audible or acoustic features to sounds themselves. Perhaps we acquire the capacity to hear subtle contrasts, pauses, and rhythms that make a phenomenological difference.

Each of the options considered above assumes that hearing speech involves hearing sounds. Speech sounds are meaningful; they are produced by voices; they have noteworthy audible qualities; but they are a type of *sound*.

Is hearing speech hearing sounds? Consider the phenomenology of speech perception. Many researchers have noted that audible speech seems, phenomenologically, like a neatly ordered, regimented sequence of distinguishable sound types known as *phonemes*, which make up distinct words organized into structured sentences. Phonemes are important in understanding the auditory perception of speech because they are the distinguishable, language-specific equivalence classes of sounds that make up the spoken words of a language. English contains approximately 40–44 phonemes, including /d/, /z/, /ʃ/ (<sh>), and /ə/. Languages differ in what sounds they distinguish as different phonemes and in what sounds they count as allophones (variants) of a single phoneme. Spanish, for instance, does not distinguish /s/ from /z/. Perhaps, then, learning a spoken language requires the capacity to hear and distinguish the phonemes that make up its words, along with their specific audible characteristics.

A great source of dispute in this area stems from a vast body of empirical research that suggests a substantial divergence between the experienced features of speech sounds and the actual features of acoustic signals. No consistent cues recognizable in an acoustic signal, such as frequency or amplitude patterns, straightforwardly determine what one hears as a given phoneme or word (see Mole, Chapter 10, and Remez and Trout, Chapter 11, for further discussion). In particular, the acoustic features that correspond to a given phoneme depend upon the phonemic context. Both prior and subsequent phonemes impact the acoustic signature for a specific phoneme occurrence. Furthermore, acoustically, different speakers differ dramatically. The main philosophical lesson is that the manifest image of speech and the scientific image of sounds appear sharply disconnected.

One response is anti-realism about speech sounds. Georges Rey (2007, 2008), for instance, argues that phonemes and other linguistic entities are mere intentional objects that commonly lack physical instances. Smith draws a similar lesson from the divergence between phenomenology and acoustics. Smith contends that while acoustic signals do not contain the linguistic sounds or structures we seem to hear, we do manage to communicate by speaking. Communication, he claims, does not require the existence of speech sounds in the world, but only requires that the world seems to contain linguistic entities.

Notice that, unless eliminativism or anti-realism is true of sounds in general, anti-realism about the objects of speech perception implies that the objects of speech perception differ from those of ordinary audition.

Some have argued explicitly that the objects of speech perception differ from those of ordinary non-linguistic audition, since, given the empirical evidence, hearing speech is not hearing sounds. For example, Liberman (see 1996) famously and influentially argues that the objects of speech perception are *intended motor commands*, since aspects of the *production* of speech, such as the articulatory gestures used to generate it, do have affinities (if not complete correspondence) with and perhaps do predict experienced phonemes. Liberman even argues that since it targets intended motor commands rather than sounds, perceiving speech invokes a dedicated perceptual module distinct from audition. Perceiving speech and perceiving sounds on this view require different perceptual modalities.

Mole (Chapter 10) is critical of Liberman's motor theory. Mole argues that it is unclear whether the motor theory is supposed to provide an account of what we experience in perceiving speech or of what is represented by subpersonal structures implicated in perceiving speech. In the former case, Mole argues, it is phenomenologically implausible. In the latter case, it is unwarranted.

Worse, it is excessively demanding and thus untenable as a claim about what is represented by the mechanisms of speech perception.

We would like to note that the empirically grounded argument that hearing speech is not hearing sounds is unsound unless sounds straightforwardly can be identified with or are determined by underlying acoustic features. There are good reasons to doubt this. Neither the sound of a car driving on a gravel road, nor the sound of wood striking wood, for example, corresponds to a simple or straightforward feature recognizable on the surface of the acoustic signal. Each is highly complex and probably requires mentioning features of its source to make its individuation intelligible. Good reasons suggest that even the qualities of pitch, timbre, and loudness lack straightforward physical correlates. Hearing speech might not be distinctive, after all.

Fowler's (1986) *direct realist* account of speech perception attempts to capture the importance of articulatory movements of the mouth and vocal tract to speech perception while arguing that speech perception is a form of ordinary, environmentally situated audition. All audition, Fowler claims, is a matter of using acoustic information to find out about things and happenings in one's environment. If ordinary audition involves awareness of sound sources, and if, counter to a very naïve physicalism, we should not expect to match features of heard sounds with straightforward acoustic features, then despite the empirical results about speech, the objects of speech perception and the objects of non-linguistic audition might belong to a common kind.

Remez and Trout (Chapter 11) draw a stronger lesson from the discoveries of speech perception science during the past century. Remez and Trout argue that no reductive account of the objects of speech perception is compatible with the empirical evidence. Thus, the objects of speech perception are reducible neither to sounds nor to intended gestures nor to articulatory movements. Instead, according to their *homeostatic properties* account, speech perception depends upon properties that are diagnostic of, but not identifiable with, particular speech sounds. The diagnostic features for speech perception might be highly theoretical and closed to introspection. To discover what such features are requires examining the processes that underlie the perception and recognition of speech.

Remez and Trout argue that the case of speech illustrates in a particularly poignant way a more general lesson. The use of introspection and phenomenological considerations in theorizing about perception and its objects requires independent justification that it has not received. They argue that without a justification, nothing of use to scientific psychology comes from examining phenomenology. For instance, they claim that considering what the experience of speech perception is like distracts from the scientific task of explaining

speech perception and linguistic understanding. Remez and Trout thus take the lessons of speech perception to warrant a general warning against relying on any methodology that uses phenomenology to discern the structure either of perceptual content or of perceptual processes.

It remains to be settled whether speech perception has special content, or has special objects other than sounds, or invokes special perceptual systems. This is fertile territory not only for conceptually sophisticated empirical work but also for philosophical and theoretical contributions. It is, however, uncontroversial that speech sounds are particularly salient and significant for humans, and that we enjoy special sensitivity to speech sounds. Human infants at a very early age distinguish speech from non-speech and show greater interest in speech sounds than in similarly complex non-speech sounds (Vouloumanos and Werker 2007). The capacities that support this interest remain to be characterized and explained.

5. Concluding Remarks

The issues we have discussed form the heart of the philosophy of sounds and auditory perception, as we understand it. The main debates concern the ontological nature of sounds; the locations of sounds; the characterization of spatial audition; whether and how we hear sound sources in addition to sounds; the nature of musical listening; and the primary objects of speech perception. We have pointed out where each question, in addition to being interesting in itself, promises to impact theorizing about perception more broadly.

We would like to conclude with a remark about a point of concern. Remez and Trout's (Chapter 11) warning about phenomenology calls into question the phenomenological constraint upon theorizing about auditory perception from Section 2.2. We claimed that phenomenology is prima-facie relevant to theorizing about the content of auditory perceptual experience. We claimed that one way to capture this is in terms of the veridicality conditions for auditory experiences, the appropriateness of which we discern in part phenomenologically. But other worries recently have been expressed about the reliability of phenomenological reports (see, e.g., Schwitzgebel 2008; Jack and Roepstorff 2003; Roepstorff and Jack 2004), and it is now commonly accepted that introspecting phenomenology is an imperfect guide to understanding perception. Phenomenological reports are influenced by numerous factors beyond just what experience presents itself as being like for

its subject. Perhaps, therefore, we should cast out introspection as a way to understand the contents and objects of audition.

We should, however, distinguish using introspection as a guide to the structure of the mind and mental processes from using introspection as a guide to how the objects of experience appear. We believe that, in the second sense, introspective phenomenology is relevant to theorizing about perception. Illusions and hallucinations are differences between how things appear and how things are. How things appear is a matter of phenomenology. How things appear impacts what we believe and what we do. What we believe and do arguably are matters impacted by perceptual content. This grounds a case for the prima-facie relevance of phenomenology to philosophical questions about the content of experience. Phenomenology in this sense also figures centrally in psychological research on perception. Theories of vision aim in part to explain why things look they way they do (see Pylyshyn 2003: ch. 1). Theories of audition similarly aim to explain why things sound the way they do. That includes explaining how audition presents things as being and why auditory experience is organized as it is (Bregman 2005). The data for these theories thus are partly introspective—they includes first-person descriptions of what can be seen or heard, and the way those things look or sound. Reports of phenomenology are data that must be explained by a psychological theory, even if only part of the explanation is that experiences have features accessible to and reportable by the subject.

What about the reliability of introspective or phenomenological methods? We should distinguish unreflective introspection from careful phenomenological inquiry. We also should not presuppose that phenomenological descriptions are obvious or self-evident. As with other data, such descriptions may be revised or rejected in the light of subsequent thinking. Done with care, however, introspection may lead to interesting insights about what perceptual experience is like and what perceptual theorizing must explain (see, e.g., the essays in Noë 2002, 2007).

Introspection, in conjunction with the kinds of philosophical methods used by the contributors to this volume, can help make clear just what any satisfactory account of auditory perception and experience must address. Many of the questions raised in the chapters that follow, such as those that concern the content of auditory experience, the experience of music, and the perception of speech, will not, however, be resolved by introspection alone. Confronting and solving these problems will require whatever insights we can glean from psychological theorizing and from philosophy. But until good

reasons show that subjectively accessible features of experiences are irrelevant to psychological theorizing—about, for instance, concepts and action—we continue to maintain the minimal thesis that, all else equal, theorizing about perceptual content should respect phenomenology.

References

Batty, C. (2007). *Lessons in Smelling: Essays on Olfactory Perception*. Ph.D. thesis. Cambridge, Mass.: MIT.

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

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

——(2005). 'Auditory Scene Analysis and the Role of Phenomenology in Experimental Psychology'. *Canadian Psychology*, 46: 32–40.

Casati, R. and Dokic, J. (1994). *La Philosopie du Son*. Nîmes: Chambon.

————(2005). 'Sounds', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2005 edition), available at <http://plato.stanford.edu/archives/fall2005/entries/sounds/>.

Crane, T. (2009). 'Intentionalism', in B. McLaughlin, A. Beckermann, and S. Walter (eds.), *Oxford Handbook of Philosophy of Mind*. Oxford: Oxford University Press.

Evans, G. (1980). 'Things Without the Mind—A Commentary upon Chapter Two of Strawson's *Individuals*', in Z. van Straaten (ed.), *Philosophical Subjects: Essays Presented to P. F. Strawson*. Oxford: Clarendon Press. Reprinted in G. Evans, *Collected Papers*. Oxford: Clarendon Press, 1985.

Fowler, C. A. (1986). 'An Event Approach to the Study of Speech Perception from a Direct-Realist Perspective'. *Journal of Phonetics*, 14: 3–28.

Gallagher, S. (2005). *How the Body Shapes the Mind*. New York: Oxford University Press.

Hamilton, A. (2007). *Aesthetics and Music*. London: Continuum.

Handel, S. (1993). *Listening: An Introduction to the Perception of Auditory Events*. Cambridge, Mass.: MIT Press.

——(1995). 'Timbre Perception and Auditory Object Identification', in B. C. Moore (ed.), *Hearing*. San Diego, Calif.: Academic Press, 425–61.

Jack, A. and Roepstorff, A. (eds.). (2003). *Trusting the Subject: The Use of Introspective Evidence in Cognitive Science*, volume 1. Charlottesville, VA: Imprint Academic.

Liberman, A. M. (1996). *Speech: A Special Code*. Cambridge, Mass.: MIT Press.

Lycan, W. (2000). 'The Slighting of Smell', in N. Bhushan and Rosenfeld (eds.), *Of Minds and Molecules: New Philosophical Perspectives on Chemistry*. Oxford: Oxford University Press, 273–89.

McAdams, S. and Bigand, E. (eds.) (1993). *Thinking in Sound: The Cognitive Psychology of Human Audition*. New York: Oxford University Press.

McDowell, J. (1998a). *Meaning, Knowledge, and Reality*. Cambridge, Mass.: Harvard University Press.

——(1998b). *Mind, Value, and Reality*. Cambridge, Mass.: Harvard University Press.

Maclachlan, D. L. C. (1989). *Philosophy of Perception*. Englewood Cliffs, NJ: Prentice Hall.

Malpas, R. M. P. (1965). 'The Location of Sound', in R. J. Butler (ed.), *Analytical Philosophy*, second series. Oxford: Basil Blackwell, 131–44.

Martin, M. G. F. (1992). 'Sight and Touch', in T. Crane (ed.), *The Contents of Experience*. Cambridge: Cambridge University Press.

——(2006). 'On Being Alienated', in T. S. Gendler and J. Hawthorne (eds.), *Perceptual Experience*. New York: Oxford University Press, 354–410.

Matthen, M. (forthcoming). 'Auditory Objects'. *Review of Philosophy and Psychology*, 7.

Neuhoff, J. (2004). *Ecological Psychoacoustics*. London: Academic Press.

Noë, A. (ed.). (2002). *Is the Visual World a Grand Illusion?* Charlottesville, VA: Imprint Academic.

——(2004). *Action in Perception*. Cambridge, Mass.: MIT Press.

——(ed.). (2007). *Phenomenology and the Cognitive Sciences* (special issue on Dennett's heterophenomenology), 6: 1–270.

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

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

——(2008). 'Object Perception: Vision and Audition'. *Philosophy Compass*, 3: 803–29.

——(forthcoming). 'Perceiving the Locations of Sounds'. *Review of Philosophy and Psychology*, 1.

O'Shaughnessy, B. (1957). 'The Location of Sound'. *Mind*, 66: 471–90.

——(1989). 'The Sense of Touch'. *Australasian Journal of Philosophy*, 69: 37–58.

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

Pylyshyn, Z. (2003). *Seeing and Visualizing*. Cambridge, Mass.: MIT Press.

Rey, G. (2007). 'Externalism and Inexistence in Early Content', in R. Schantz (ed.), *Prospects for Meaning*, volume 3 of *Current Issues in Theoretical Philosophy*. New York: de Gruyter.

——(2008). 'In Defense of Folieism: Replies to Critics'. *Croatian Journal of Philosophy*, 8: 177–202.

Roepstorff, A. and Jack, A. (eds.) (2004). *Trusting the Subject: The Use of Introspective Evidence in Cognitive Science*, volume 2. Charlottesville, VA: Imprint Academic.

Scholl, B. J. (2001). 'Objects and Attention: The State of the Art'. *Cognition*, 80: 1–46.

Schwitzgebel, E. (2008). 'The Unreliability of Naive Introspection'. *Philosophical Review*, 117: 245–73.

Scott, M. (2001). 'Tactual Perception'. *Australasian Journal of Philosophy*, 79: 149–60.

Scruton, R. (1997). *The Aesthetics of Music*. Oxford: Oxford University Press.

Smith, A. D. (2002). *The Problem of Perception*. Cambridge, Mass.: Harvard University Press.

Smith, B. C. (ed.) (2007). *Questions of Taste: The Philosophy of Wine*. Oxford: Oxford University Press.

Strawson, P. F. (1959). *Individuals*. New York: Routledge.

Tye, M. (2000). *Consciousness, Color, and Content*. Cambridge, Mass.: MIT Press.

Vouloumanos, A. and Werker, J. F. (2007). 'Listening to Language at Birth: Evidence for a Bias for Speech in Neonates'. *Developmental Science*, 10: 159–64.