Hearing Where Things Are

Hearing Where Things Are

By

D. L. Oxtoby, B.A. (Hons.)

A Thesis

Submitted to the School of Graduate Studies

in Partial Fulfilment of the Requirements

for the Degree

Master of Philosophy

McMaster University

McMaster University MASTER OF ARTS (2013) Hamilton, Ontario (Philosophy)

TITLE: Hearing Where Things Are

AUTHOR: D. L. Oxtoby, B.A. (Hons.)

SUPERVISOR: Dr. Sandra Lapointe

NUMBER OF PAGES: viii, 124

ABSTRACT

One of the central questions in the philosophy of sounds and hearing is the question of space: what spaces or locations, if any, do sound perceptions make one aware of? When I hear a sound, do I perceive the direction of the sound? The direction (or distance) of the sound's source? The boundaries or dimensions of the space the sound is produced in, or of the source itself? And if sound perceptions do make one aware of space, then with what level of determinacy?

In the first chapter of this essay, I describe my approach to sounds and hearing, and state what I take to be the fundamental challenges for any view of sound perception. For one, I take the everyday experience of sounds to be one of the most significant obstacles to an account of sound perception, and one that has scarcely been recognized as such. In everyday hearing, we are not the least bit concerned with sounds. We use sounds to gather information about the behaviour of their sources, which are typically the object of our attention whenever we perceive a sound. If I hear the sound of a car honking or a person speaking, I immediately pay attention to the car and how I can avoid it, or to the person and the meaning they intend to communicate. In everyday hearing, our awareness of sounds is similar to our awareness of windowpanes while watching the goings on outside. Consequently, the everyday experience of sounds is problematic as a model of sound perception.

In the second and third chapters, I discuss the two most popular views of sound perception in the philosophical literature, the remote view and the non-spatial view. Since these views have received much attention in the literature, I spend more time raising objections to them in chapter III than describing them in chapter II. One of the principle aims of this essay is to make the case that both of these views are mistaken, despite the valuable insights contained in each.

In the fourth and fifth chapters, I discuss the medial view. While the idea that sounds are sound waves located in a medium is the predominant view of sounds themselves in auditory science and the history of philosophy, the view that we hear sounds to be located in the medium has received little attention. Some objections to the medial view have been raised, which I address in chapter V, but very little has been said to defend or even describe the medial view. Part of the motivation for this essay is that I am struck by the fact that the medial view, which would seem to follow naturally from auditory science and the history of philosophy, has been so little discussed. Consequently, the bulk of this essay is dedicated to a description and defence of the medial view.

ACKNOWLEDGEMENTS

Special thanks go to my adviser, Dr. Sandra Lapointe, whose advice and dedication has greatly increased the quality of this thesis. Your comments on my work, financial support for research activities, and guidance in the world of professional philosophy have been invaluable. Thank you!

Dr. Casey O'Callaghan and Dr. Mark Johnstone gave the thesis a careful reading and have made several helpful comments for which I thank them. I would especially like to thank Dr. Casey O'Callaghan for the time he took to meet with me and for his very helpful feedback on my work. Thanks also to Dr. Charles Siewert for taking the time to read and discuss excerpts of this thesis.

A big thank you to my colleagues Dan Coren, Adam Sopuck, Noel Iverson, Josh and Joanna Zaslow, and Michael Hemmingsen for reading and discussing my work, and especially for your friendships these past two years. All of my friends and colleagues at McMaster have made my time here an absolute pleasure.

I would also like to thank the McMaster Philosophy department for their financial and moral support. One can only hope for such a superb environment to work in.

Finally and most importantly, I would like to thank my family for their love and support. Very special thanks go to my parents Bill Oxtoby and Bonnie Walker, to my step dad Bob Walker, and to my grandparents Don and Betty Kibbler. One does not come so far without such a wonderful family.

Contents

I

Sounds and Hearing

1.1: *The Philosophical Problem*

The primary aim of this essay is to understand how we hear where things are. This question is made more interesting by the fact that the scientific study of sounds and hearing has been one of the fastest growing areas of biomedical research in the last thirty years.[1] In parallel to auditory science, the philosophical study of sounds and hearing is gaining attention in the philosophy of perception, having received only a cursory treatment in the historical literature. With attention to the recent explosion of auditory science, the philosophy of sounds and hearing offers not only an up to date theoretical framework for understanding sound perception, but an additional model to that of vision for understanding sense perception in general.[2]

It is easy to see why there is a scientific problem about hearing where things are. The ear is an intricate mechanism which is still not fully understood. Moving from the ear to the brain, the picture only becomes more complicated, and sounds themselves have also proven difficult to model in physics and acoustics.[3] But why is there also a philosophical problem about hearing where things are?

The answer given by Malpas, 1965 is that even a fully detailed scientific account of

1 Moore, David R. "Series Preface". *The Oxford Handbook of Auditory Science*. Ed. Christopher J. Plack. Vol. 3. New York: OUP, 2010. Page *v*.

2 O'Callaghan, 2007 discusses the importance of alternative models to that of vision for understanding perception (pages 1-12).

3 The attempt to give a scientific account of pitch provides a dramatic example. See De Cheveigné, 2010.

the conditions under which we hear where things are would not fully explain the role played by conscious experience.[4] Malpas brings out the difference between the conditions for hearing where things are and the conscious experience of hearing where things are by way of an example: suppose a prospector is exceptionally good at finding underground deposits of uranium without the aid of any instruments. Scientists could explain the conditions under which the prospector can tell where the uranium is located, which might include a high dosage of radiation that the prospector is subject to when in close proximity of a deposit. But, Malpas argues, this would leave unanswered the question of the sense in which the prospector is aware of conditions like the dosage of radiation; “the problem is unsolved because we have not been told how we come to be aware of these conditions.”[5]

I take Malpas to be saying that there is a philosophical problem about hearing where things are because the scientific study of sounds and hearing describes the conditions under which we hear where things are, but the conscious experience of hearing where things are is not itself one of these conditions. That is, the experience of hearing where things are is distinct from the conditions which cause the experience, like sound waves striking the ear, or the firing of the auditory nerve. On Malpas' view, this can be seen from the fact that the prospector, if asked how he can tell where the uranium is, could not explain the experience of locating the uranium by reporting the existence of certain conditions, like that of a high dosage of radiation. Absent any instruments for detecting radiation, the prospector is oblivious to high levels of radiation, just as hearers are

4 Malpas, 1965. Page 133.
5 Ibid. Page 134.

oblivious to the mechanics of sound waves or the functioning of the ears. For Malpas, the philosophical problem is to give an account of what the prospector is aware of that would explain how the prospector finds uranium deposits. So to in the case of sounds: “We cannot then ignore the question of what we are aware of when we locate sounds, for it is what we are then aware of that enables us to understand the meaning of sentences such as 'The sound is coming from over there'”.[6]

I think Malpas' characterization of the philosophical problem of hearing where things are is fundamentally correct, but we can get a clearer idea of the problem by building on his characterization. Ultimately, the philosophical problem is not to understand the meaning of sentences about sounds or to understand what we are aware of when we locate sounds, if the latter is taken to mean the “object” or cause of our experience.[7] We want to know about the awareness itself; what is the nature of our experience of hearing where things are and how does this experience function in the rest of our mental life, such as in forming beliefs about where things are. The philosophical problem is to account for the nature and function of the experience of hearing where things are, while the scientific problem is to account for the conditions under which we have the experience.

6 Ibid. Page 136.

7 Perceptions are usually said to have an “object” in two difference senses, the “material sense” and the “intentional sense”. The material object of a perception is the feature of the external world which causes a perceptual awareness of that object. For example, one might take sound waves to be the material object of a sound perception. The intentional object of a perception, on the other hand, is whatever the subject is aware of, whether it be a material object that one is perceptually aware of or some feature of one's pereptual experience which one simply attends to. For example, the intentional object of a sound perception might be sound waves, or it might be the pitch or the loudness of sound waves (Anscombe, 1965). In this essay, I do not use the intentional sense.

Chapter I

1.2: *Sound Experiences and Sound Perceptions*

How we hear where things are is a large question in the contemporary philosophical literature on sound perception and is connected to the even larger question of how we perceive locations or spatial relations at all. The focus of this essay is on one part of the problem of hearing where things are, namely the epistemology and phenomenology of sound perceptions. I call the epistemology and phenomenology “one part” of the problem because both are concerned with the experience one has in hearing where something is rather than the nature of sounds themselves as a feature of the external world.

Sound perceptions, like all perceptions, put a subject into relation with something in the world. The subject has an experience, and something in the world answers to or is the “object” of that experience, be it a material object, a state of affairs, or perhaps an event. The epistemology and phenomenology of sound perception is focused on this experience rather than its object. What knowledge is made available in the experience of a sound, and what is the nature of this experience? Suppose I hear a sound coming from my open window. Do I merely know the pitch, loudness, and timbre of the sound? Do I also know where the sound itself is located, or even further, where the sounding object is located? And what are the features of the sound experience in virtue of which I have this knowledge?

The experience of a sound, however, cannot be studied in complete isolation from its object, sounds themselves. One important feature of the experience of a sound is the extent to which it is accurate; that is, whether and with what level of detail the experience is a *bona fide* perception rather than an illusion. The question of illusion is also strongly connected to the epistemology of sound perceptions: if a sound experience is supposed to be a fully fledged perception, then the hearer must be aware of some properties that the corresponding sound actually has, like its location, pitch, loudness, and so forth. This requires a comparison between the features of a sound experience and the features of the corresponding sound itself. If I experience a sound coming from my open window as coming “from the left”, what sort of thing would the sound itself have to be in order for this experience to be a fully fledged perception? In other words, what would have to be true of the sound itself in order for it to really be coming from the left?

I take sound perceptions to be a subset of sound experiences, distinguished by the fact that perceptions are experiences that make the hearer aware of sounds as they really are. The expression 'sound experience', on the other hand, I use to cover cases where I wish to consider a sound experience without considering whether the experience is veridical. For example, suppose one has the experience of a sound coming from the left. If there actually is a sound coming from the left, then the experience is a perception. If not, then the experience is an illusion. Often, however, I will only wish to talk about the experience without considering whether it is an illusion or a perception, and here I will use the term 'sound experience' rather than 'perception' or 'illusion'.

1.3: *Hearing Where “Things” Are: Sounds and Sounding Objects*

Another question to distinguish from those of auditory experience and of sounds themselves is that of the “source” of a sound. In the philosophical literature, the word 'source' is usually meant to refer to a sounding object (while it is sounding).[8] Examples of sounding objects include a running car, a speaker that is sounding, or a bird that is chirping. There is an ambiguity, however, as sometimes the word 'source' is used to refer to merely the location of a sounding object, as in Malpas, 1965.[9] In this essay, I will always use the word 'source' to refer to the causal origin of a sound, which I take to be a sounding object. If one hears the sound of a trumpet, the sound is one thing and the trumpet is another. I call the trumpet the “source” of the sound.[10]

The important difference between sounds and sources is that we do not perceive sources through hearing alone. If I hear the sound of a car, the sound is the object of my perception and the car is something I may or may not infer the presence of. If I am trying to work inside my office, I will hear the *sound* of cars on a busy street outside, but I will not be distracted by any thoughts about the cars themselves. I need not infer (or deny) the presence of a source just because I hear its sound. The presence of the sound, however, is something I am directly aware of in the sense of having perceptual access.

Through correlation with visual and tactile experience in early childhood development, one learns very fast that the things one sees and feels are the sources of sounds. We even apply this knowledge to our future perceptions, and come to expect that certain sounds mean a certain source (“sounds like a storm is coming!”). But one does not have perceptual access to the source or anything about it (a point I return to and substantiate in

8 I use “object” in the widest possible sense. A lightning bolt, for example, could be the source of a sound. A harder case is that of sounds produced simply by changes of pressure in the air. I hesitate to call a change of pressure an “object”, so this may be a case of a sound without a source (though, of course, not without an origin).

9 Malpas, 1965. Page 136.

10 This is not to say that sounds might not, in some way, be properties of their objects (see Pasnau, 1999; Kulvicki, 2008). But properties are not identical with their objects. Even if sounds are properties of their objects, one could distinguish between the object (the source) and one or more of its properties (its sound).

sections 4.2 to 4.5 and 4.9). For example, one could not tell *just by hearing the sound* whether or not its source is a real storm or a set of high quality speakers playing an engineered sound. Moreover, if hearing were the only sense one ever possessed, the thought of sounding objects would never arise.[11] From this we should conclude that sounds are we have perceptual access to on the strength of hearing alone, and never their source.

An important consequence of the assumption that sounds are the exclusive object of hearing is connected to the expression “hearing where *things* are”. Usually when one says that one hears where something is, one is referring to a sounding object, rather than the sound itself. We tend to speak of hearing where the car is or where the bird is, rather than where the sound of the car or the sound of the bird is. But if sounds rather than sounding objects are what one hears, how can one learn the location of sounding objects just by hearing sounds?

In the philosophical literature, there are two ways of responding to this question. One response is that we do not hear where sounding objects are just by hearing sounds. Sound perceptions alone do not make the hearer aware of any location whatsoever (the “non-spatial” view of sound perception).[12] The other response is that sound perceptions do indeed make the hearer aware of locations (“spatial” views of sound perception, which come in three varieties). However, since sounds and not their sources are the objects of hearing, the location one is perceptually aware of, on spatial views, is the location of the sound. How one gets from the perceived location of the sound to the location of the

11 Strawson, 1959 makes a similar point (page 78).

12 Reid, 1764; Strawson, 1959; O'Shaughnessy, 2000; Nudds, 2009; Scruton, 2009 subscribe to versions of this view.

source depends on which of the three spatial views one adopts. The following chapters of this essay are dedicated to describing and criticizing both spatial and non-spatial views.

1.4: *Sounds in Everyday Hearing*

The claim that sounds are the only objects of hearing should not be thought to conflict with the common circumstance where hearing a sound is part of a larger experience in which one otherwise perceives the source (probably by seeing it). If I hear the sound of a car and I also see the car, then I perceive both the sound and its source. But, while the source is at least one of the objects of my perception, it is not the object of my auditory perception, which is just to say that I do not have perceptual access to the source by hearing its sound. As far as I can tell from the sound alone, the source could be a car, a set of speakers, or someone who can do a good vocal impression.

I emphasize this point because it is common to say things like “I heard the car” rather than “I heard the sound of the car”. The former expression suggests that the car itself is the object of a sound perception, and such expressions are of common use. But if sounds are the exclusive objects of sound perceptions, why do we commonly speak as if their sources are also objects of hearing?

I think the answer to this question is simple, but it leads to a deeper point. We commonly speak as if sounding objects are what we hear because sounding objects, in day to day life, are what we are interested in. If I hear the sound of a car coming toward me, I do not pause to contemplate its sound. I use the sound to get the information I need to avoid the oncoming car. Similarly, if someone is talking to me, I do not pause to appreciate the sound of their voice. I might notice it, but if I were to attend to the sound of a person's voice, I would be distracted from the message they intend to communicate. In fact, if I hear a sound and I cannot immediately call to mind what sort of thing is making it, I feel a bit disturbed and say “what was that?”

In everyday hearing, we hardly pay any attention to sounds. We care about sounding objects or what sounds signify to us, and we perceive sounds like we perceive windowpanes. Certainly, they are something we perceive, but usually our attention is almost exclusively elsewhere. We attend to what is outside the window, and what is making the sound, and only occasionally to the windowpane itself or the sound itself. The deeper point is that, if we are hardly aware of sounds in everyday hearing because we rapidly attend to their source or their imagined significance, then everyday hearing is unlikely to tell us much about the experience of sounds or of hearing where sounds are.

1.5: *The Problem of Sound Individuation*

Everyone capable of hearing sounds will know intimately what the experience of a sound is, though a detailed account of this experience is much harder to give than it may seem at first. Sounds themselves, on the other hand, are much more mysterious. Even if one knows what it is like to hear a sound, though this experience may be difficult to describe in detail, one still has only the slightest idea of what in the external world a sound is. Are sounds whatever causes us to have an auditory experience? If so, we have a causal chain starting with sounding objects and their properties causing sound waves, sound waves causing reverberations in an acoustic environment, collections of sound waves and reverberations causing ear drum vibrations, and so on – we are left with little idea as to which part of the causal chain is a sound. I will call this the “problem of sound individuation”.

Some philosophers have recently attempted to solve the problem of individuation by defining sounds as whichever link in the causal chain best accounts for the experience of hearing a sound.[13] For example, if systems of sound waves can be shown to bring about all or most of the features of an auditory experience in a way that gives rise to little or no illusion, or in any case less illusion than competing candidates for what a sound is, then these philosophers would conclude that systems of sound waves are the best candidate for sounds. In general, there is consensus in the philosophical literature that, all else being equal, the less illusion a theory of sounds attributes to the hearer, the better the theory. That is not to say that a theory which attributes even constant illusion, an “error theory”, is necessarily false. But a theory of equal explanatory power that avoids making hearing into a systematic illusion is always preferable to one that does not.

I go into some detail about the individuation problem in sections 5.4 and 5.5, but my focus in this essay is the experience of hearing where things are. I focus on the experience of hearing where things are because an account of this experience can act as a basis for solving the individuation problem: each link in the causal chain extending from sounding objects to the hearer can be thought of as occupying a discreet location; sounding objects are in one place, the medium which contains collections of sound waves is in another place, and the hearer is in yet another place. Any answer to the individuation problem must locate sounds somewhere along this causal chain, but if the answer is to locate

13 O'Callaghan, 2007 and Casati and Dokic, 2010 are examples.

sounds counter to where we experience them to be, then this answer to the individuation problem produces an error theory at least with respect to hearing where sounds are, and should therefore be held in suspicion.

1.6: *The Problem of Auditory Experience*

The question of where we experience sounds to be located is usually thought to have one answer and to apply to all cases: we either experience sounds to be where sounding objects are, *or* in the medium, *or* where the hearer is. These three candidates exhaust the full range of spatial views; that is, the full range of any view of sound perception which holds that, in a sound perception, the hearer experiences some spatial features of the perceived sound, like its direction. It might also be suggested that the view that sounds are heard to be in the medium could also be combined with the view that sounds are heard to be at the source, saying that sounds are heard to be near the source or in the neighbourhood of the source. I raise some criticisms of this “combined view” in section 4.8.

The problem of auditory experience is, where do we experience sounds to be located? This problem is connected to the problem of where sounds are in fact located, but the two should not be confused. The problem of auditory experience is about the nature of an experience, whereas the problem of locating and defining sounds themselves is about a feature of the external world.

For each of the three locations that we might experience sounds to be located at, there is a corresponding theory of sounds themselves which locates them there. Some philosophers say that sounds are properties of sounding objects[14] or events[15], either of which should be thought of as something located exactly where the source is. Most scientists yet fewer philosophers say that sounds should be thought of as sound waves[16] or collections of sound waves[17] located in the medium. Still fewer philosophers say that sounds should be thought of as bodily phenomena or sensations produced by bodily phenomena[18] located where the hearer is.

On the view of sound perception that I defend in this essay, the “medial view”, hearers experience sounds to be located in the medium which surrounds the hearer and the source of a sound. Normally, the medium is the air surrounding a hearer and the sounding objects in their environment, but anything which is conducive to sound waves could, in principle, be a medium. For example, if one hears a sound underwater, then the water surrounding the hearer also acts as a medium. If hearers do in fact experience sounds to be in the medium, then sounds themselves would have to be something in the medium, like sound waves or collections of sound waves, in order for sounds to be perceived without any illusion about their location. I state and defend the medial view in chapters IV and V respectively.

A more widely endorsed view of sound perception, the “remote view”, holds that sounds are heard to be located exactly where the source is.[19] In order to avoid attributing any illusion to hearers about where sounds are located, proponents of the remote view must identify sounds with something located exactly where the source is, like properties

14 Pasnau, 1999; Kulvicki, 2008.
15 Casati and Dokic, 2005; 2010; O'Callaghan, 2007; Pasnau, 2009.
16 O'Shaughnessy, 1984.
17 Perkins, 1983; Sorensen, 2007; Nudds, 2009.
18 Evans, 1980; Maclachlan, 1989.
19 Perkins, 1983; Casati and Dokic, 2005; 2010; Pasnau, 1999; O'Callaghan, 2007; 2010. Kulvicki, 2008.

of the source or events located exactly where the source is. I describe the remote view in chapter II and raise some objections in chapter III.

The “non-spatial view” of sound perception is an interesting and widely held alternative to spatial views. According to the non-spatial view, we do not hear sounds to be located anywhere at all.[20] Sounds, as features of the external world, certainly have a location, but sound perceptions simply do not have the capacity to make the hearer aware of a sound's location. By correlating the sounds one hears with other information, however, hearers are able to infer the location of sounds. For example, if one hears a sound and can see its source or can remember where its source would be located given that it is making such and such a sound, then one can infer where
the sound is located. Since the non-spatial view holds that a perceived location is absent from sound perceptions, it cannot be an error theory of sound perception when it comes to hearing where sounds are; sound perceptions cannot be systematically in error about the location of sounds if they do not present any location whatsoever. I describe the non-spatial view in chapter II and raise some objections in chapter III.

The view that sounds are heard to be located where the hearer is enjoys little if any support. Even those who hold that sounds themselves are in some sense located where the hearer is do not claim that they seem to be so.[21] Consequently, I focus on the above three view, which I now turn to.

20 Strawson, 1959; O'Shaughnessy, 2000; Nudds, 2009; Scruton, 2009.
21 O'Shaughnessy, 2000; Nudds, 2009.

II

The Remote View and the Non-spatial View

2.1: *The Remote View of Sound Perception*

It is sometimes claimed that the remote view of sound perception is the natural, "untutored" description of everyday hearing.[22] Some opponents of the remote view, including this account, grant that the remote view has some intuitive appeal. Pasnau, 1999, Kulvicki, 2008, and O'Callaghan, 2010 suggest that the remote view's appeal lies in the fact that it can explain the basic function of hearing in everyday life, navigating one's environment.[23] Frequently, hearers use sounds to gather information about the location of sounding objects in their environment. According to the remote view, hearers get information about where the sounding object is located by perceiving sounds to be located in the exact same place as the sounding object.

But how could hearers get from the perception of a sound to the awareness of something distinct from the sound, namely the location of its source? Whatever one is aware of in hearing arises from perceiving a sound, so we are left with a question as to how information about the source's location enters into the perception of a sound.

The remote view sets itself apart from other spatial views in two ways: (i) by its

22 Pasnau, 1999. Page 317. O'Callaghan, 2010. Page 133. 'Untutored' appears in Casati and Dokic, 2010. Section 1.2.

23 Pasnau, 1999. Pages 317-318. Kulvicki, 2008. Page 9. O'Callaghan, 2010. Pages 123-124.

account of what spatial information sound perceptions have and (ii) how that information enables the hearer to locate sounding objects. On the remote view, sounds are perceived to be located in exactly the same place as their source, making the hearer aware of the location of the source simply in virtue of making the hearer aware of the location of the sound. The location of the source is identical to the perceived location of the sound, so perceiving the location of a sound is sufficient for perceiving the location of its source. The perceived location of the sound of a car honking from two lanes over, for example, is such that the hearer is perceptually aware of a sound which is two lanes over in exactly the same place as the car. Since "two lanes over" is not only the perceived location of the sound but also that of its source, the hearer has a perception of the source's location *in virtue* of having a perception of the sound's location.[24]

2.2: *The Epistemological Argument for the Remote View*

The remote view's claim that the perceived location of a sound is identical to the location of the source can be supported by an epistemological argument. If sounds are not perceived to have the same location as their source, how could hearing alone generate knowledge of the location of the source? In other words, if hearing alone does not present hearers with the location of a sound's source, how could hearing alone be used to navigate one's environment?[25] Since hearing alone can surely be used to navigate one's environment (though it might not be easy), it seems that hearing alone must present the

24 The idea that sounds are the immediate object of hearing and sounding objects are perceived by way of their sounds is standard among the views discussed, but I go into more detail about this in sections 1.3, 4.2 to 4.5, and 4.9.

25 O'Callaghan, 2010. Page 123.

hearer with the location of sounding objects. The argument could be stated as follows:

1) Hearing a sound makes the hearer perceptually aware of the location of that sound.
2) In order to learn where a sound's source is located by hearing alone, the perceived location of a sound must be identical to the location of the source.
3) Hearers are readily able to learn the location of the source by hearing alone.

∴ 4) Hearing a sound makes the hearer perceptually aware of the location of the source.

The first premise is a description of sound perception that is not a point of contention among spatial views, but is rejected, along with the other two premises and the conclusion, by the non-spatial view. The idea that hearing the location of a sound is part of hearing a sound, however, warrants some explanation.

Suppose one hears a sound *s* and suppose that *s* is located at a place *p*. Spatial views hold that hearers learn what location *p* is simply by hearing *s*, which can be understood as a claim about the structure of sound perceptions: all sound perceptions have the form "*s* is at *p*".[26] Hearers learn what location *p* is because a perception of *p* is always part of a sound perception. The level of detail with which hearing can make one perceptually aware of locations is uncertain, but if a hearer has the perception "*s* is at *p*", then *p* is at least perceived as "wherever *s* is located", and *p* is in this sense known to the hearer. The upshot is that because perceiving the location of a sound is part of perceiving a sound, sound perceptions are sufficient for knowledge of where the sound is located.

The second premise in the above argument contains a claim that differentiates the

26 Where *s* is understood to also have certain auditory qualities (volume, pitch, timbre, rhythm, etc.).

remote view from other spatial views, but before considering this claim it is worth noting that the argument is deductively valid. The third premise is merely a description of hearers' everyday experience of hearing that only non-spatial views will reject, so if the first two premises are accepted in conjunction with the account of everyday hearing in the third premise, then the conclusion follows necessarily.

The defining claim of the remote view in the second premise connects the perceived location of sounds to the location of the source. One, it seems sure, is able to learn the location of a sound's source by hearing alone, for it would be possible to navigate one's environment by hearing alone if one had to. It would therefore seem that hearing alone must furnish the hearer with knowledge of the source's location. Yet, by hypothesis, the only location sound perceptions present the hearer with is that of the sound, so how does information about the source's location get into the perception of its sound?

Supposing that p is the perceived location of a sound and q is the location of its source, the claim contained in the second premise is that the only way a sound perception of the form "s is at p" could furnish the hearer with knowledge of q is if, for all sound perceptions, $p = q$. If p and q are identical, then a sound perception "s is at p" is sufficient for a perception of q (since it is sufficient for a perception of p). On the other hand, if p and q are not identical, then sound perceptions would not be sufficient for a perception of q and hearers would not necessarily learn where q is in virtue of perceiving p. But hearers seem well able to locate the source of a sound by hearing alone so, according to the remote view, we should accept the second premise. In order to learn where a sound's source is located, sound perceptions must be such that the perceived location of the sound is identical to the location of the source.

In virtue of what properties are sounds perceived to be located at their source? Malpas, 1965 and O'Shaughnessy, 1965 describe two such properties, the perceived direction and perceived distance of a sound.[27] Firstly, when one hears a sound, it seems to come from a certain direction. Hearers can, by hearing alone, readily tell the difference in direction between a call for the hearer's attention from the left and a call for the hearer's attention from the right. On the remote view, the perceived direction of a sound must therefore be the direction of the source (this follows from the remote view's claim that the perceived location of the sound is identical to the location of its source).

In addition to a perceived direction, sounds also seem to come from a certain distance. Hearers can tell the difference between hearing the sound of a car crash from two hundred feet away and hearing the sound of the same car crash from just twenty feet away. But how to understand perceived distance is, as with perceived direction, a major point of contention among spatial views. On the remote view, the perceived distance of a sound is the distance of the source (which also follows, as it does in the case of perceived direction, from the remote view's claim that the perceived location of the sound is identical to the location of the source). The sound of a car crash two hundred feet away is heard to be fixed in place, like the crash itself, at a distance of two hundred feet away.

2.3: *Sounds, Colours, and Qualities of Objects*

One motivation for the remote view extends the analogy between sounds and colours. One might be inclined to hold that sounds are heard to be in the medium because, without

27 Malpas, 1965. Page 137. O'Shaughnessy, 1964. Page 77.

a medium, one would be unable to hear a sound.[28] It is also true, however, that without the medium of light, one cannot see colours. Yet few philosophers are willing to concluding from the fact that one cannot see colours in the dark that colours are "in the light". While a medium may be required to perceive sounds, it does not follow necessarily that sounds are in the medium; after all, the fact that light is required to perceive colours is not normally taken to force the conclusion that colours are in the light (though Russell, 1912 and O'Shaughnessy, 1984 subscribe to versions of this view). Rather, one might hold with Pasnau, 1999 that sounds are like colours: sounds are properties of sounding objects located exactly where the sounding object is.[29]

But perhaps sounds are not analogous to colours. If sounds and colours are both located exactly where the sounding or coloured object is, how is it that sounds can be perceived without the immediate presence of the sounding object, whereas colours are always seen in conjunction with a coloured object? The sound of a person speaking can be heard from another room, from around the corner, etc., but the colour of the person's shirt cannot be seen independently of seeing the shirt. Pasnau, 1999 calls colour perception "direct" for this reason (whereas sound perception is indirect), but concludes that "of course, this is not a satisfactory basis for saying that the one exists in the object, the other in the air. The difference seems merely a consequence of the different physical properties of light *versus* sound waves. Whereas light waves cast shadows, sound waves can pass around objects, and can bounce off common objects with far less distortion."[30] The "directness" or "indirectness" of sound perceptions, Pasnau argues, should not be

28 Pasnau, 1999. Pages 322-323. For discussion, see Pasnau, 2009. Pages 38-40. O'Callaghan, 2007. Pages 47-55.
29 Pasnau, 1999. Page 322.
30 Ibid. Page 323.

thought of as a determining factor in where sounds are heard to be located.

Pasnau, 1999 also defends the remote view's account of perceived location from a possible objection. One reason that sounds might seem to be located elsewhere than their source is that sounds do not behave like their source or its properties in two important respects. Unlike most sounding objects and most of their properties, like their colour, shape, etc., sounds are quite short lived and are only heard when an object is modified in a way that produces vibrations, as when objects collide. If sounds were a property of their source located where the source is, it seems intuitive that sounds should behave like other properties of their source, which are relatively stable over a lengthy period of time and require no modification of the object.[31]

For example, the colour of one's furniture, the scent of a piece of cheese, and the size of one's car remain, for the most part, fixed as they are until something modifies them (the cheese biodegrades, the car is totalled). Sounds, on the other hand, are fleeting and dependent on the modification of an object whose other properties seem fixed and stable. The speaker is 1' by 4', 5 lbs., and green in colour, but only sounds when its cones are made to vibrate. As a result, the speaker seems to have a certain size, weight, and colour, but it does not seem to “have” a sound – it merely *makes* a sound when it is modified by an electrical current. It seems that the sound must be something separate from the speaker, which might lead one to conclude that sounds are heard to be something separate from their source.

Pasnau responds with a thought experiment:

31 Ibid. Page 322.

Chapter II

> In a world full of objects that make a constant, continuous noise as a matter of their own intrinsic nature (they need not, for instance, be struck), we would say that some objects have a noise, others do not, and we would classify objects in terms of the character of their noise (squeakers, murmurers, etc.). . . But if sounds are in objects in that world, then should they not also be in objects in this world? The fact that sounds here are typically short-lived should not make a difference to where they are located.[32]

The thrust of Pasnau's thought experiment is that the perceived duration of a sound has nothing to do with its location. If the world had turned out such that many objects were constantly sounding, there would be no hesitation to locate the sound within its source. The belief that sounds are separate from their source simply arises from the short lived nature of sounds, but, according to Pasnau, beliefs about sounds, their properties, or their sources should not get in the way of characterizing sound perception. We should set aside these beliefs and see that sound perception is more accurately described accord to the remote view, namely as locating sounds at the distance and direction of the sounding object, just like one sees colours, shapes, and sizes to be located exactly where their object is.

An insight in Pasnau, 1999 is that features of sounds and hearing that might seem peripheral to the perceived location of sounds can influence ones account of perceived location. Perceived duration, it seems to Pasnau, has nothing to do with perceived location, yet the transient nature of sounds suggests that they are unlike properties such as colour or shape, and therefore that sounds are perhaps located elsewhere than colour and shape. One might also be taken aback by the fact that sounds can be heard without the immediate presence of the source, which seems to suggest that sounds are somehow

32 Ibid.

separate from their source. But these facts, according to Pasnau, do not determine where sounds are heard to be located. In section 4.9, I will return to these seemingly peripheral features of sounds and hearing, arguing instead that the capacity to perceive sounds without the immediate presence of the source and the perceived duration of sounds are important desiderata for an account of the perceived location of sounds.

2.4: *The Non-spatial View of Sound Perception*

According to the non-spatial view of sound perception, sounds are not heard to be located anywhere at all. When one hears the sound of a car horn while driving along a busy street, the sound of the horn, strictly speaking, is not heard to have a location. One may hear the sound of the horn and know where it is coming from, but one does not know where the sound is coming from by hearing alone. Sound perceptions contain no spatial information and are merely correlated with other information available to the hearer that genuinely is spatial, like that which is made available in seeing or remembering where a sound is coming from.

It is not a contentious claim that the non-spatial view does not reflect the everyday experience of hearing. In normal cases, when one hears a sound, one attends to its source and immediately knows the location of the source. However, by calling attention to some interesting challenges and possible counterexamples to spatial views, the non-spatial view presents reasons for thinking that the everyday experience of hearing may conceal the nature of sound perceptions.

For example, one could imagine waking up to the sound of loud music that seems to be coming from a party somewhere down one's street. By hearing alone, however, one could easily be mistaken about exactly where the sound is coming from. Suppose that, upon further inspection, the sound turns out to be much quieter music coming from a roommate in another room, or even quieter music coming from trick speakers hidden around one's room. In each case, what one takes to be a loud sound coming from far away turns out to be a quieter sound coming from a closer location.

In this example, the hearer's perception of the music (and especially the perceived volume of the music) is the basis for the hearer's beliefs about where the sound is located. The information presented in hearing the music, however, can be seen to be insufficient to accurately inform the hearer of the sound's location. This suggests that sound perceptions, while perhaps providing cues that are correlated with the location of the sound, do not actually present the location of the sound. On the non-spatial view, hearers may use cues like the perceived volume of a sound to surmise or infer its location, and these cues may be further correlated with memories about where similar sounds have been located or a visual perception of the source of the sound in order to produce more reliable beliefs about the sound's location, but nevertheless, sound perceptions themselves contain no genuinely spatial information.

The idea that hearers do not perceive the location of sounds but surmise it from correlated information is also illustrated by the engineering of sounds in films. If an event like an explosion or a car crash is portrayed at a distance in the background of a scene, the accompanying sound is quieter than if the crash had been in the foreground of the scene. Yet all of the sounds in a film come from the same place, namely the speakers making them. In order to manipulate the experience of “distance”, sound engineers manipulate correlated auditory qualities like volume, reverberation, and so forth.

In another example, if one were blindfolded and unaware that there was a soundtrack of mundane sounds being played on speakers placed around a room, like the sounds of a kettle whistling and cars going by outside, one would not judge all of the sounds to be located inside the room. While the sound of the kettle might be located with relative accuracy, the sound of the cars really would seem to be coming from outside. Assuming good quality, well placed speakers, hearers would be helpless to tell where the sound of the cars is coming from without removing their blindfold.

The converse happens in "sound externalization" experiments. Hearers wear a pair of headphones which play sounds that are engineered to seem as if they are coming from various locations around the hearer. Even though the source of the sound is obviously the headphones, subjects report that the sound so strongly seems to be coming from the location it is engineered to suggest that the sound from the headphones is indistinguishable from the real thing.[33] Such cases can motivate the non-spatial view's claim that whatever idea hearers have of a sound's location, it is not based on directly perceiving that location (after all, subjects in sound externalization experiments do not perceive the sounds to be coming from their headphones). Because the perceived location of a sound varies by varying other auditory qualities, like volume, reverberation, and especially the time of onset in the case of externalization experiments, perceived location may seem to be mediated by other qualities rather than directly perceived. Pitch, on the other hand, does not vary by the volume or time of onset of a sound, and therefore does not seem to be inferred from other qualities.

33 Hartmann and Wittenberg.1996. Pages 3678-3688.

O'Shaughnessy, 2009 suggests that, in addition to other qualities like volume or reverberation, hearers' beliefs about the location of sounds vary strongly by their background cognitive states and abilities:

> For example, one hears a knock on the door, and it can easily seem to one that the knock is an audible entity situated over there by the door. But now suppose that one were to close one's eyes, and suppose also that one had been set down in some grossly unfamiliar, constantly shifting scene—where then would the sound appear to be? Now make the sound an utterly unfamiliar noise. What is left of spatial perception?. . . do not these considerations strongly suggest that one's *mind* sets the knock at the door...[34]

How familiar the hearer is with the sound and the environment has a strong effect on where a sound is believed to be coming from. If a hearer knows that a sound *s* coming from a place *p* will be perceived to have *xyz* auditory qualities, then this background knowledge will lead the hearer to believe that for any *s*(*xyz*), *s*(*xyz*) is located at *p*. The sound perception itself, however, consists only of perceived auditory qualities *xyz* and does not include a perception of *p*. Whereas spatial views hold that sound perceptions can be said to have the form "*s* is at *p*", the non-spatial view holds that sound perceptions have the form "*s* has *xyz*", and therefore do not make the hearer perceptually aware of any place *p*. By learning what auditory qualities sounds have when they are coming from various places, hearers acquire the background cognitive ability to associate sounds with places and thereby become "familiar" with certain sounds and certain environments.

2.5: *Strawson and the Non-spatial View*

The idea that spatial information is absent from sound perceptions takes the form of a

34 O'Shaughnessy, 2009. Page 124.

general theory in Strawson, 1959. Strawson invites the reader to imagine a purely "auditory world" that is entirely composed of sounds with all of their auditory qualities of volume, pitch, timbre, etc.[35] Strawson notes that one difficulty of the experiment is that, in a world of only sounds, there is no room for the physical body of the perceiver, but nevertheless concludes that auditory qualities alone are insufficient for establishing in a perceiver any concept of space. Consequently, sound perceptions contain no spatial information. The difficulty posed by the hearer's body is that the composition and arrangement of the ears on the head is enough to hear a sound as either "to the left" or "to the right" of one's body, but if we eliminate this difficulty along with any access to the other senses, then sounds would be perceived only in relation to other sounds. One could only describe sounds with words like "louder", "softer", "sharper", "faster", " more or less dissonant", " more or less consonant", etc., and none of the relations expressed by such terms gives sufficient information to establish anything about the location of a sound.

Visible objects, in contrast, always have spatial relations to each other. If a blue and a red patch are perceived, the blue patch must be "beside", "below", or in any case in a "different place" than the red patch, else the blue patch occludes or blends in with the red. Since visible objects are perceived to have a location relative to each other, perceiving two visible objects is sufficient to establish in the perceiver the notion of a spatial relation ("to the left", "to the right", etc.) In the real world, sounds must of course have some kind of spatial relation to each other if they are to be in the external world at all, but the case of an auditory world is meant to suggest that if sounds are understood to be bundles of

35 Strawson, 1959. Pages 65-66.

auditory qualities like volume, pitch, and timbre, then sounds cannot be perceived to have any spatial relations, and therefore cannot be said to contain any spatial information that the hearer may have perceptual access to.[36]

If sounds do not contain any spatial information, sounds must be assigned a location after they are perceived by correlating the perceived auditory qualities with other qualities that genuinely do contain spatial information.[37] For example, when I hear the car two lanes over honking, I know it is two lanes over because I may simultaneously see the driver honking, or I may remember that the last time I heard just *that* sound, I saw a driver honking from two lanes over. The hearer correlates the sound of the honk with the visible or remembered location of the source, but the hearer does not directly hear the location of the sound or its source.

2.6: *Correlation and the Association of Sounds with Locations*

An account of how hearers associate certain sounds with certain locations is an important feature of non-spatial views. If sounds are not consciously perceived to have a location or to have the location of the source, how does hearing allow one to navigate one's environment? I have suggested a few possibilities, which I will now state in the form of three models of the association of sounds with locations.

The first model is that sounds may be associated with a location by way of correlated information from other sense modalities. One may associate sounds with a certain location because one sees or otherwise perceives the location of the sound's source, as in

36 Malpas, 1965 offers a similar interpretation of Strawson's thought experiment (page 140).
37 Strawson, 1959. Page 63.

the case where I associate the sound of the car horn with the car two lanes over because I see the driver honking. Associating sounds with locations by way of the other senses is the most basic and direct model of how hearers may come to beliefs about the location of a sound in the absence of any spatial information from hearing alone, “basic” in the sense that it provides a mechanism by which spatial information is presented to the hearer and “direct” in the sense that this information is directly presented in perceptions from the other senses.

A second model of the association of sounds with locations is by memory. One may associate a sound with a certain location because one remembers that just such a sound was, in the past, located in just such a location, as in the case of the soundtrack: the hearer will associate “outside” with the soundtrack of cars driving by in the distance because this is how the hearer remembers cars driving by in the distance to sound. If, in the first instance, hearers associate sounds with locations by correlated information from the other senses, then it is easy to see how association by memory gets its start.

A third model of association could easily be produced by a conjunction of the previous two over some period of time. As O'Shaughnessy suggests, sounds may be associated with locations by way of a background ability that one acquires by compounding correlated information from the other senses and memories of where sounds have been located in the past (Reid thought this is how we develop a fully fledged sense of hearing[38]). The hearer has so much practice at using correlated information from the other senses and is able to remember and predict what sounds will sound like in a given location so well that they can be said to have acquired a background ability for locating

38 Reid, 1764. Pages 28-29.

sounds. When a sound is heard, its location is “automatically” known to the hearer as a result of this background ability.

These three models of associating sounds with location could be interpreted as models of what hearers consciously do to learn the location of a sound, at least until one acquires the background ability to “automatically” associate sounds with locations. This is one way to understand the association of sounds with places, but non-spatial views may instead hold that association is entirely an unconscious or sub-personal cognitive process. One seems to hear sounds and know where they are coming from all at once as the result of what is sometimes called an “unconscious inference”.

In the philosophy of perception, the notion of unconscious inference traces back to the natural philosopher Alhazen[39] (circa. 11th cent.), and receives its modern formulation in Helmholtz's optics.[40] According to Helmholtz, unconscious inference occurs when an incoming sound produces a physiological response in the hearer that the auditory system recognizes and matches to a corresponding idea about the sound. The metaphor is that the auditory system is “inferring” from the physiological reaction caused by a sound to certain ideas or cognitive states like a sound perception or a belief about where the sound is located.

Non-spatial views of auditory experience might try to make use of the notion of unconscious inference to explain hearing where things are. When one hears a sound, one does not hear it to have any spatial information whatsoever. Instead, ones hears auditory qualities like pitch, timbre, and loudness, and these qualities trigger in the hearer

39 Hatfield, 2002.
40 Helmholtz, 1867. Pages 28-29. For discussion see Hatfield, 2002. De Cheveigné, 2010. Pages 84-87.

unconscious inferences that are responsible for the hearer's beliefs about where the sound is coming from. While an unconscious inference about where a sound is coming from might happen too fast for the hearer to notice, non-spatial views maintain that, nevertheless, incoming sounds do not contain any spatial information.

On the other hand, if all of the unconscious processes are in place to allow the hearer to associate sounds with locations, it does not seem unlikely that these same processes could also make the hearer perceptually aware of a sound's location.[41] If a hearer has acquired a background ability for associating sounds with locations, then it could be argued that the hearer possesses the apparatus necessary for consciously perceiving the location of a sound. Proponents of the non-spatial view, however, may concede this point and merely insist that hearing, taken by itself, is not responsible for making the hearer perceptually aware of a sound's location. If a hearer must acquire a background ability for locating sounds by first associating sounds with locations via the other senses or via memory, then hearing alone is not responsible for locational information. Consequently, it might be argued, sound perceptions, taken by themselves, do not contain locational information.

The remote view and the non-spatial view are the two most discussed views in the philosophical literature. Whereas the remote view holds that sound perceptions have the form "*s* is at *p*" where $p=q$ (the location of the source), the non-spatial view holds that sound perceptions merely have the form "*s* has *xyz*" and are absent any spatial information. I think that both of these views are mistaken. Having outlined the remote

41 This point was suggested to me in a discussion with Professor O'Callaghan. See also O'Callaghan, 2007. Pages 40-45.

view and the non-spatial view along with some of their motivations, I now offer some objections to each view.

III

Objections to the Remote View and the Non-spatial View

In the last chapter, I talked about three views of sound perception and went into some detail about the two views that are most discussed in the current philosophical literature, the remote view and the non-spatial view. In this chapter, I present five objections to the remote view, followed by a refutation of the non-spatial view on both theoretical and practical grounds. For reasons I hope to make clear, both views must be rejected in favor the medial view, which I state in chapter IV and defend in chapter V.

3.1: *Objections to the Remote View of Sound Perception*

A key motivation of the remote view is the claim that it can explain how hearers learn the location of sounding objects without committing one to a view of sound perception that is inconsistent with our experience. There are cases, however, that present problems for both the phenomenology and epistemology of the remote view's account of perceived location. In these cases, the remote view both misdescribes the experience of a sound's location and overestimates the knowledge of a sounding object's location that can be gleaned from sound perceptions. As a consequence of misdescribing the phenomenology in such cases (discussed in sections 3.2 and 3.3), the remote view is forced to attribute auditory illusion about the perceived location of the sound.

But illusions about direction and distance are not the only forms of illusion that arise on the remote view. The remote view's account of perceived loudness and perceived acoustic qualities attribute pervasive forms of illusion (discussed in sections 3.4 and 3.5). The following four sections go into detail about these four forms of illusion.

In section 3.6, a more general objection to the remote view's account of perceived location is raised. I argue that the remote view's account of perceived location depends upon the implausible notion that hearing makes one aware of the boundaries of sounds and their sources. I call this the “boundary problem”.

Some proponents of the remote view claim that it should be counted in the remote view's favor that it attributes so little illusion to sound perception.[42] However, if the remote view is forced to attribute at least the following four forms of illusion, then the remote view makes hearing into a highly illusory affair relative to the medial view.[43] The addition of the boundary problem to these four forms of illusion makes the remote view hard to maintain in light of the alternatives.

3.2: *Perceived Direction*

On the remote view, illusions about the direction of a sound arise in cases where there is an obstruction which obscures the path of sound waves travelling form the source to the hearer. The remote view holds that the perceived location of a sound is that of the source, and therefore that the perceived direction of a sound is the direction of the source. In

42 Pasnau, 1999. Pages 314-316. Casati and Dokic, 2010. Section 1.2. O'Callaghan, 2007. Page 29.
43 I will return to the question of the medial view and attributions of illusion in chapter V.

cases of an obstruction between the hearer and the source, however, sounds do not seem to lie in the direction of their source. If a hearer is standing in a quiet room with an open window at some distance from the

hearer, all of the noise from outside seems to lie in the direction of the open window regardless of the direction of sounding objects outside. If a car is driving by outside, the sound of the car always seems to be coming from the open window regardless of the location of the car.

This experience is common to anyone living in a noisy city and is relatively widespread, since some degree of obstruction is a common feature of one's environment. Sounds from outside seem to come from the direction of open windows or doors, sounds from around corners seem to come from the direction of the corner itself, and sounds from the far side of obstructions like buildings seem to come from above or beside the obstruction (especially if the hearer is standing near the obstruction). Such illusions are frequent enough that, in the presence of obstructions, hearers become accustomed to ignoring the perceived direction of the sound as a basis for learning where its source is located, or sometimes know how to compensate in order to accurately infer the direction of the source. Of course, not all obstructions are misleading about the direction of the source. If I hold up a pencil between myself and a conversation partner, the sound of their voice does not seem to have a direction different than that of its source. Cases of large scale obstruction, however, are common enough to conclude that the remote view attributes moderately widespread illusion to our capacity to perceive the direction of a sound.

Such cases show that sound perceptions do not make the hearer perceptually aware of the direction of the source. In the case of hearing sounds from outdoors, sounds seem to lie in the direction of the open window between the hearer and the source. To hold, with the remote view, that sound perceptions make the hearer aware of the direction of the source would yield the false conclusion that the source lies in the direction of the open window, since this is the direction that sounds from outdoors seem to come from.

If possible, it is preferable to avoid attributing illusion to the hearer. One might ask, following

the lead of Austin's discussion of perceptual illusion,[44] what is the illusion that one is supposed to be subject to here? Take the case of sounds coming from outdoors. The sounds from outdoors seem to come from an open window quite independently of the location of their source. But, while the hearer is likely to guess wrongly about the direction of the source, there is something the hearer would certainly not guess wrongly about, and which it would hardly make sense to guess about in the first place, namely that there is a sound coming from the direction of the open window. The hearer is under no illusion when it comes to the direction of the sound, though the perceived direction of the sound would lead the hearer to *infer* wrongly that the source lies in the direction of the open window.

In this case, the perceived direction is the direction of the open window, so if something is perceived without illusion, it must be something that lies in the direction of the open window. But all that lies in this region of space are sound waves in the air. sound waves are certainly a candidate for sounds themselves, and one that would be consistent with the medial view of sound perception, but the remote view cannot identify

44 Austin, 1962.

sounds with something located elsewhere than the source without attributing pervasive illusion. On the medial view, rather than attributing illusion in cases of obstruction, one can instead conclude that one has no illusion about the direction of the sound, which is heard to lie in a region of open space in the medium extending from the hearer to the open window, but that the hearer might wrongly *infer* the direction of the source.

The objection from perceived direction can also be formulated as a phenomenological argument or an epistemological argument against the remote view. The phenomenological argument is that the remote view fails to accurately describe sound perceptions when there is an obstruction between the hearer and the source. Imagine that the sound of a bird chirping is

coming from the direction of an open window on the far side of a room, and there is a closed window right in front of the hearer through which the bird can be seen. In such a case, one can see the source of the sound straight ahead, yet the sound seems to be coming from the direction of the open window on the far side of the room. Since the remote view is committed to the claim that the perceived direction of the sound is the same as the direction of the source, the remote view is committed to a description of sound perception that does not match up with our experience in such cases.

The epistemological argument is that the remote view implies that sound perceptions make more knowledge about the direction of a sound's source perceptually available than they actually do. If the perceived direction of a sound is the same as the direction of the source, then one can learn the direction of the source simply by perceiving the direction of the sound. In the above cases, however, hearers do not learn the location of the source by hearing alone. Take the case of the chirping bird. If the hearer were blindfolded and oblivious to their surroundings, the hearer would wrongly judge that the bird really does lie in the direction of the open window.

3.3: *Perceived Distance*

Another phenomenon that the remote view cannot accommodate is engineered sounds, specifically those which are engineered to mislead the hearer about the distance of the source. Since engineered sounds are not always designed to suggest that the source is somewhere it is not, and since one's day to day experience may not be rife with engineered sounds, this form of auditory illusion is not pervasive. Nevertheless, an explanation of problematic cases of engineered sounds that does not attribute illusion to the hearer is preferable to one that does.

Sounds in a theatre are a good example. According to the remote view, the perceived distance of a sound is the same as the distance of the source, but if this were true then the experience of
going to the theatre, as we know it, would be impossible. Consider the perceptual situation of a subject watching a film at a theatre. The subject hears many sounds from the film which create the illusion that the source of the sound is quite far away, like the sound of a distant explosion in the film, and others which create the illusion that the source of the sound is quite nearby, like the sound of quiet dialogue between characters in the foreground of a scene. The subject's auditory experience can be described as follows: the subject has the experience that some sounds are coming from far away, like an explosion portrayed as far away in the film, and that some sounds are coming from nearby, like

dialogue in the foreground of the scene. This description is reinforced by the likelihood that, if the subject were blindfolded and oblivious to the fact that they were at a theatre, the subject would believe that some sounds actually were coming from far away and others from nearby (assuming that the speakers had good sound quality).

On the remote view, however, sound perceptions present sounds to be located at the same place as the source, and it is no mystery that the source of all of the film's sounds is the theatre's speakers. In a theatre, the location of the sounds' source is always the same, yet one's experience of a film's soundtrack is always such that the distance of a sound seems to match whatever distance serves the purpose of the film rather than the distance of the theatre's speakers. If an explosion is portrayed as happening two hundred metres away relative to the lay of the scene, then the explosion really sounds like it is coming from two hundred metres away. If it did not, then the sound of the explosion would not match the explosion portrayed on the screen.

In general, all of a film's sounds must be experienced to be coming from a distance that matches the distance of the source portrayed in the film. If the apparent distance of the sounds does not match the portrayed distance of the source and instead seems to have the distance of the real source, the theatre's speakers, then all of a film's sounds will seem to be coming from a distance limited to the confines of the theatre, making our experience of the film totally
unrealistic. An explosion portrayed as two hundred metres away could never seem to come from further away than the theatre's speakers.

Again we might ask, what is the nature of the illusion the hearer is supposed to be subject to? In the case of the theatre, clearly one does not hear the sounds to be located at the theatre's speakers, as the remote view would predict. Instead, hearers are misled about the location of a sound's source because sound engineers produce sound waves that mimic the properties that sound waves would have if their source really was at the distance of the events depicted in the film. The sound waves corresponding to the sound of the approaching car mimic the properties that sound waves would have if they were actually produced by an approaching car, which involves engineering the intensity, reverberation, clarity, and other acoustic qualities.[45]

The fact that holding a sound's source constant while varying the sound's qualities can mislead one about the distance of the source suggests that the distance of the source is arrived at by inference rather than perception. If the distance of the source were perceived along with the location of the sound, then varying the above sound qualities should not mislead one about the distance of the source. Since this is clearly not the case with certain engineered sounds like those at the theatre, it would seem that sound perceptions end at the location and qualities of the sound and do not extend to the distance of the source, which can only be inferred. A sound's location and qualities are typically a good guide to the distance of the source because a sound's qualities are usually determined by the source in a highly regular way. Quiet, reverberant sounds are usually made by sources that are far away and loud, clear sounds are usually made by sources that are nearby. But when these qualities are engineered, if one tries to determine the distance of the source by way of its sound, one may wrongly *infer* the location of the source.

If there is good reason to accept that the location of the source is inferred rather than perceived, as the medial view holds, then the hearer is not subject to illusions about

45 Ibid.

perceived location in the case of engineered sounds. The sound's qualities are accurately perceived; the sound of a car driving toward the hearer really does get louder, really does decrease in reverberation, really does become clearer, and so forth. All of these auditory qualities have been engineered, but they are all accurately perceived by the hearer. If the medial view is correct, the sound is perceived to be in the medium between the source and the hearer, though one may wrongly infer the distance of the source. Taking the theatre example, the engineered sounds are heard to be in the medium in the direction of the theatre's speakers, but the qualities of the sound may easily mislead the hearer about the distance of the speakers. Whereas the sound is perceived without illusion, the hearer may wrongly infer the distance of the source.[46]

As with the objection from perceived direction, the objection from perceived distance can be formulated as a phenomenological argument or an epistemological argument. The phenomenological argument is that the remote view fails to accurately describe the experience of engineered sounds. In the case of going to the theatre, one experiences sounds that are engineered to suggest that their source is nearer or further away depending on the portrayed distance of what is being shown in the film. If a film's sounds were not perceived to have the portrayed distance of events in the film but to have the distance of their source, the theatre's speakers, then the film would seem totally unrealistic. A vehicle approaching from two hundred metres away, for example, would always seem to have the distance of the theatre's speakers, which never changes, or at any rate is never two hundred metres. In such cases, sounds are not experienced to have the distance of their

46 I discuss how hearers *rightly* infer the distance of the source according to the medial view in section 4.2 and 4.4.

source.

The epistemological argument is that the remote view implies hearers learn the distance of a sound's source by hearing alone in cases where they actually do not. If sound perceptions present

sounds to have the same distance as their source, then hearers may learn the distance of the source just by hearing its sound. This is false, however, in the case of sounds that are engineered to suggest that there source is nearer or further away than it actually is. In the case of visiting a theatre, if a hearer is blindfolded and oblivious to the fact that they are in a theatre, then they would judge incorrectly that the source of the film's sounds is at whatever distance the sound was engineered to suggest to the hearer. On the screen, a car approaching from two hundred kilometres away really does sound like a car approaching from two hundred kilometres away. Assuming the speakers are of sufficient quality, the hearer, by hearing alone, would judge incorrectly that the source of the sound is at the distance that the sound is engineered to suggest rather than that of the actual source (the speakers).

3.4: *Perceived Loudness*

On the remote view, perceived loudness is a pervasive form of illusion. The apparent loudness of a sound, we know from everyday hearing, depends on how close the hearer is to the source. The further sound waves travel from the source of a sound to the hearer, the more energy is lost and the quieter the sound becomes. Since the remote view claims that sounds are not located in the medium, this variability in apparent loudness is a distortion of the actual loudness of the sound (i.e. the loudness of the sound at its source). Since sound waves must travel in order to exist, any sound waves that reach a hearer will distort the loudness of the sound at the source, making every instance of perceived loudness, on the remote view, illusory.[47]

It is worth considering Pasnau, 1999's argument for the remote view's conception of perceived loudness:

> The standard view [the wave theory of sound] can offer a seemingly natural explanation [of perceived loudness]: sounds seem more intense when one is closer to the origin of the sound because they *are* more intense. I must disagree. On my view, sounds only seem more intense when one is closer to their source. In fact the sound may not have changed in intensity; all that may have changed is one's sensation. One has the sensation of a louder sound, because one is closer to the sound. . . My proposal may strike the reader as implausible. But our practice of measuring sounds in terms of the intensity of sound waves is not as natural as one might suppose. On this system we have no straightforward way of measuring how much sound an object makes. We find it natural to ask how loud a jet engine is, or how loud a concert is, but all we can say in answer to such questions, on the standard view, is that the engine or the concert has such and such a sound when measured from a certain distance. Answers of the latter form will often be precisely what is wanted, and will perhaps always be good enough. But still there seems something peculiar about the standard view's inability to answer a simple question: how much noise does it make?[48]

While Pasnau is addressing the wave theory of sound rather than the medial view of sound perception, the above passage is relevant to the medial view since sound waves or something else in the medium must be the (material) object of sound perceptions on the medial view. The argument in this passage, which Pasnau calls the “measurement objection” to the wave theory of sound, is that despite the apparent implausibility of the remote view's account of loudness, any view which locates sounds in the medium will

47 Pasnau, 1999. Pages 319-320.
48 Ibid.

also be unable to give a plausible definition of loudness. Such a definition would have to be able to answer the supposedly simple question, how much noise does a sounding object make? If the measurement objection goes through, then it appears there would be no reason to favour the remote view over the medial view when it comes to perceived loudness. The measurement objection, however, can be addressed in the following way.

Firstly, it is not true that, strictly speaking, all we can say about loudness is that sounds have a certain loudness at a certain distance. In recent philosophical literature on the wave theory, how much noise an object makes has been described as the resulting sound waves' level of pressure.[49] This is certainly also the mainstream scientific view, which accepts the wave theory and describes loudness in terms of the pressure of waves in the medium: “Because loudness [perceived loudness] is primarily correlated with sound intensity, loudness is most often displayed or described as a function of physical sound intensity or pressure”.[50]

Taking sounds as bodies of sound waves and loudness as their level of pressure, it can be seen that the loudness or pressure of a sound can be calculated at any given point in the body of sound waves, or it can be calculated for the body of sound waves as a whole. Since the pressure at any given point is determined by the displacement of air at that point,[51] the “total pressure” of the body of sound waves can be construed as an additive function of the displacement of air at each point within the body of sound waves. If the question is how loud a sound seems to the hearer, then the relevant pressure is that of the sound waves at the location of the hearer. But, if the question is how loud a sound is in

49 Perkins, 1983. Page 176.
50 Epstein and Marozeau, 2010. Page 46. See also Davies, 2010. Page 375.
51 See Nave, 2013 for a discussion of pressure and air displacement. <http://hyperphysics.phy-astr.gsu.edu/hbase/waves/standw.html#c2>

total (i.e. Pasnau's question, “how much noise does it make?”), then the relevant pressure is the total pressure of the body of sound waves. Total pressure has no significance to perception since sounds are always perceived from a particular location, but if it is necessary to give an objective measurement of how much noise something makes (in total), there is in principle no reason that this could not be measured as an additive function of the loudness at any given point.

Also relevant is Pasnau's suggestion that a description of sounds in terms of their loudness at a particular distance is usually the form of description one wants anyway, and that such descriptions might always be good enough. I think all parties agree to this point. Since the *perceived* loudness of a sound always depends on the distance of the hearer from the source, descriptions of how loud a sound is heard to be always presuppose that the hearer is at some distance or another. If someone asks me how loud a concert was, they want to know how loud I heard it to be, which implicitly depends on where I was standing. They are usually not interested in the total amount of air displaced by the body of sound waves that the musicians produce. Unless an auditory scientist is asking me how loud the concert was and really does want an objective answer, there is no problem. But even in such a case, auditory scientists measure perceived loudness primarily in terms of the pressure of a body of sound waves, so I could in principle answer the auditory scientist about the “total loudness” of the concert even in this special case.

The measurement objection, while raising interesting questions about how to measure loudness, is not ultimately a problem for the wave theory of sound, and therefore not a potential problem for the medial view of sound perception. Since the quantity of pressure of sound waves can be measured either at the location of the hearer or in total, the wave theory has no problem about measuring the loudness of sounds. Moreover, because the remote view holds that sounds merely *seem* to vary in loudness by the hearer's distance from the source, the remote view makes every experience of loudness illusory, since the hearer is always at some distance or another from the source.

3.5: *Perceived Acoustic Qualities*

Another form of pervasive illusion that follows from the remote view is that of perceived acoustic qualities. Sounds seem different depending on the nature of the environment they are produced in. The sound of a violin seems different in a small room with hard surfaces than it does outdoors. What accounts for this difference is the interaction of sound waves with the environment they are produced in, the result of which is the perceived acoustic qualities of a sound. Perceived loudness and reverberation are obvious examples; in a small room with hard surfaces, a violin sounds louder and much more reverberant than it does outdoors because sound waves immediately reflect off of the surfaces and lose very little energy depending on the hardness of the surface. Since the sound waves need not travel far and are not dampened by soft surfaces, very little energy is lost, resulting in a greater perceived loudness than would be experienced outdoors. And since the sound waves are not dampened by soft surfaces, much of the sound waves are preserved as they are reflected at the hearer, creating the experience of reverberation or “echo”. The scientific literature tends to agree on adding two more acoustic qualities to this list, perceived clarity and spaciousness.[52] Without going into detail, a sound is said to seem “clear” to the extent that sound waves reach the hearer

52 Davies, 2010. Pages 299-300.

directly (without being reflected off of surfaces in the environment), and “spacious” to the extent that the sound waves arriving at each ear are dissimilar as a result of lateral reflections off of surfaces in the environment.

Perceived acoustic qualities must be illusory on the remote view for the same reason as perceived loudness, considered as it is in the above section as the output loudness of the source. The reason perceived loudness is illusory on the remote view is that one's perception of loudness is caused by sound waves, which, according to the remote view, do not present sounds as they are at the source. Since sound waves lose energy as they travel, sounds seem quieter as one moves away from the source. But sound waves are also modified by the environment they are produced in, resulting in a sound's perceived acoustic qualities. If sounds are located at the source, then they do not reflect off of the surfaces of the environment, and therefore do not have any degree of reverberation or any other acoustic qualities. Perceived reverberation must therefore be illusory, and so to with the rest of the acoustic qualities since all perceived acoustic qualities are produced by the modification of sound waves by their environment.

3.6: *The Boundary Problem*

Since the remote view holds that sounds are heard to have the location of their source, the remote view must also be committed to the claim that sounds are heard to have boundaries. The argument can be formulated in two versions, one weaker and one stronger. The weaker version of the boundary argument is directed at the remote view in general, whereas the stronger version is directed specifically at the notion of perceived distance. The weaker version can be stated as follows:

1) The remote view holds that sounds are heard to be located exactly where the source is.

2) The remote view holds that hearers are perceptually aware of the location of the sound/source, since this is how one learns the location of the source when one hears its sound.

3) The location of the source typically has rigid boundaries, namely the boundaries of the source itself. If a car is the source of a sound, then the boundaries of the source's location are the boundaries (or "edges") of the the car.[53]

∴ 4) The remote view is committed to the claim that hearers are perceptually aware of rigidly bound locations, since the source of a sound typically has a rigidly bound location.

In this version of the boundary argument, the remote view is said to be committed to the claim that sound perceptions make the hearer perceptually aware of rigidly bound locations. If sounds share the location of the source, and the location of the source is defined by the rigid boundaries of the source itself, then the location of the sound also has rigid boundaries that the hearer is said to be perceptually aware of.

The claim that hearers are perceptually aware of the boundaries of a source is one natural way to interpret the remote view. It is obviously false, however, that hearing makes one perceptually aware of a source's boundaries. If it did, the blind should have a

53 One might think that the car is not really the source of the sound; the engine, or some part of the engine is the source. This does not challenge the boundary argument, since any part of the car that is specified as the source will also have rigid boundaries. The problem merely shifts from one set of boundaries to another.

much easier time getting around. The fact that hearing does not make one perceptually aware of a source's boundaries can also be seen in one's own experience. If I close my eyes and listen to the sound of a car, I am not perceptually aware of the boundaries of the car.[54] I could not trace a diagram of the car's boundaries on the strength of hearing alone, and I will not be perceptually aware of the car's boundaries until I open my eyes or reach out and touch the car.

That hearing does not make one perceptually aware of a source's boundaries is also reflected in the fact that one cannot identify the shape of a source by hearing its sound. If I close my eyes and am presented with a sound resembling that of a car, I might open my eyes and be surprised to find a stack of speakers in front of me. In such a case, if I were perceptually aware of the boundaries of the source, then I would have known that the source did not have the shape of a car. Hearing makes no such information available, unlike vision or touch.

This version of the boundary argument is "weak" in the sense that the remote view is easily reformulated to avoid the claim that sounds are heard to have rigid boundaries. Rather than claiming that sounds are heard to have the same rigidly bound location as their source, a proponent of the remote view might claim that sounds are merely heard to be located at some point within the boundaries of the source's location. On this reformulation of the remote view, if I hear the sound of a car, the sound seems to be located at some point within the boundaries of the car's location (probably the approximate centre of the car, or the approximate centre of the engine if we take the engine to be the source of the sound). The point at which the sound is heard

54 See footnote 53.

to be located might be though of as the point of the greatest concentration or "intensity" of sound.[55] The car and the engine have rigid boundaries, but their sound is not heard to share the boundaries of its source. Rather, the sound seems to be concentrated at some point within the boundaries of the source.

The stronger version of the boundary argument challenges this reformulation of the remote view. It can be stated as follows:

1) If sounds are heard to be at the location of the source and not to share the boundaries of the source, but to be concentrated at some point within those boundaries, then sounds are still perceived to have a distance from the hearer (namely, the approximate distance of the source).

2) The remote view holds that hearers are perceptually aware of the distance of the sound/source, since this is how hearers learn the distance of the source.

3) Perceived distance implies perceived boundaries. To perceive a sound to lie at one distance and not another implies that one is perceptually aware of the distance at which the medium, which is soundless on the remote view, gives way to the sound.

∴ 4) The remote view is committed to the claim that hearers are perceptually aware of the boundary between the medium and a sound, since the distance of this boundary from the hearer is the distance of the sound from the hearer.

The crucial premise of this version of the boundary argument is that perceived distance implies perceived boundaries. Even if the remote view is reformulated to hold that sounds are perceived to be a some point of concentration within the boundaries of the source,

55 I further develop this idea throughout chapter IV in the context of the medial view.

sounds must nevertheless, on the remote view, be perceived to have some distance. If sounds are not perceived to have some distance, and in particular the distance of their source, then sounds cannot be perceived to be located where the source is. Consequently, the remote view's claim that hearers learn the location of the source in perceiving the location of its sound must be false.

The boundary argument, however, cannot be used to challenge the reformulated version of the remote view with respect to perceived direction. If one merely perceives a sound's point of greatest concentration to lie in a particular direction, one need not perceive the sound to have any boundaries. A sound's lying in a particular direction does not require that the sound be separated from the hearer by a stretch of the medium (which is soundless on the remote view). One merely hears a sound, wherever its boundaries might be, to lie in a certain direction. Perceived distance, on the other hand, depends on being perceptually aware of the distance at which the medium gives way to a sound. It is this stretch of the medium leading up to a sound that makes a sound “distant” from the hearer, and it is the boundary where this stretch ends that marks the particular distance of a sound. If sounds are not heard to have a boundary where the medium gives way to a sound, then sounds are not heard to be at a remote distance from the hearer, and the remote view must be false.

The force of the boundary argument is that no matter how one reformulates the remote view, proponents are committed to the idea that hearing makes one perceptually aware of boundaries. The remote view is committed to the claim that either sounds are perceived to share the boundaries of the source or to share a boundary with the medium at a remote distance from the hearer. Both of these claims seem implausible.[56] Upon hearing a sound, one is hard pressed to identify the boundaries of the source or the boundaries of the sound. Hearing seems to operate with insufficient determinacy to establish boundaries, unlike vision or touch.

3.7: *The Case Against the Remote View*

While the remote view has a coherent explanation of how hearers learn the location of sounding objects, the counterexamples of hearing in the presence of obstructions and hearing engineered sounds (sections 3.2 and 3.3) show that there are cases which the remote view cannot accommodate. In the case of hearing in the presence of obstructions, the perceived direction of the sound is not that of the source; in the case of engineered sounds, the perceived distance of the sound need not be that of the source. Since these counterexamples attack both the remote view's account of perceived direction and perceived distance, it appears that neither part of the remote view can be salvaged. Moreover, in these cases, the remote view gets both the phenomenology and epistemology wrong. The perceived location of the sound does not seem to be that of the source, nor does one learn the location of the source in perceiving its sound.

The main appeal of the remote view is that it provides an intuitive explanation of how we navigate our environment by hearing: we know the location of sounding objects because sound perceptions make the hearer perceptually aware of sounds at just that location. It is a consequence of this claim, however, that cases where sounds are misleading about the location of their source become cases of auditory illusion on the

56 Sorensen, 2007, for example, denies that sounds are heard to have boundaries (page 284).

remote view. The only way to avoid attributing illusion in such cases is to hold that sounds can be heard to be located elsewhere than the source, and therefore to reject the remote view. The case against the remote view that emerges is therefore two pronged: on the one hand, the boundary problem presents a theoretical objection to the remote view, and on the other, cases where sounds are misleading about the location of the source present specific counterexamples to the remote view's account of perceived location that cannot be accommodated without attributing auditory illusion.

The above counterexamples also force us to reconsider what location, if any, sound perceptions make the hearer aware of. If sounds help one learn the location of the source, it must be by coincidence or correlation, which seems to open the door for the non-spatial view. If hearers learn the location of sounding objects by correlating their sound information from the other senses or from memory, then the further claim that hearers learn the location of sounds themselves by correlation begins to appear plausible. In the following section, I argue that it is not, and therefore that the non-spatial view does not accurately describe sound perceptions.

3.8: *Objections to the Non-spatial View of Sound Perception*

According to the non-spatial view, sound perceptions do not make hearers perceptually aware of the location of sounds. One perceives certain auditory qualities like pitch, timbre, and loudness, but one does not perceive a location. Motivations for the non-spatial view can be lumped into two categories. One might be skeptical about the perceived location of sounds out of theoretical considerations, many of which trace back to

Strawson, 1959.[57] But one might also be skeptical about perceived location out of practical considerations arising from cases where hearers seem unable to tell where sounds are coming from, some examples of which I will discuss. I argue that both the theoretical and practical considerations that might be mustered in support of the non-spatial view fail to establish the conclusion that sound perceptions do not present the location of sounds.

3.9: *Strawson and Theoretical Considerations for the Non-Spatial View*

According to Strawson, 1959, whatever it is about sounds that gives the hearer the idea that a perceived sound has a location, it is not hearing alone.[58] Strawson examines the nature of sounds at length, going into detail about pitch, timbre, loudness, duration, and rhythm (but especially pitch). After a discussion of what a perceiver could derive from these qualities, Strawson concludes that sounds cannot not provide the necessary resources to produce in perceivers any concept of space. There is nothing about the pitch, timbre, loudness, or any other qualities of sounds that could ground any concept of spatial location whatsoever, neither perceived direction nor distance.[59]

When it comes to *sounds themselves*, much of what Strawson has to say is as interesting as it is, in my opinion, correct. His account of sound qualities, the individuation of sounds, and the nature of sounds is rich and detailed in a way that I have not endeavoured to describe. The reason for my brevity, however, is that his argument

57 Strawson, 1959 is cited as the theoretical basis for the non-spatial view in Malpas, 1965. Pages 140-142. O'Callaghan, 2007. Page 29. Casati and Dokic, 2009. Page 101.
58 Strawson, 1959. Page 66.
59 Ibid.

about *sound perceptions* is invalid. One could agree with everything Strawson has to say about sounds themselves and still reject his conclusion about sound perceptions, namely that sound perceptions alone do not make the hearer aware of the location of sounds.[60]

The reason for thinking Strawson's argument is invalid is that whatever one says about sounds themselves need not be true of sound perceptions. It does not follow from the claim that sounds themselves do not contain the resources to produce an idea of location that sound perceptions do not contain the resources to produce an idea of location. Sound perceptions depend on more than just the audible qualities of sounds; they also depend on the body of the hearer and whatever mental state the hearer must be in for the effect of a sound on the hearer's body to produce a sound perception. It would therefore be a fallacy of composition to draw conclusions about sound perceptions merely from premises about sounds themselves. Sound perceptions are neither identical nor similar to sounds themselves, so what is true of sounds themselves need not be true of sound perceptions.

Strawson intentionally omits the role of the hearer's body:

> Before leaving the auditory world altogether, I should consider a possible objection to the whole procedure of this chapter. . . I selected the model of the auditory world as one from which [physical] bodies were altogether absent. . . But by what right do I assume the possibility of such types of experience, and of such schemes? . . . They are models against which to test and strengthen our own reflective understanding of our own conceptual structure. Thus we may *suppose* such-and-such conditions; we may discuss what conceptual possibilities and requirements they can be seen by us as creating. . . In all this we need no more claim to be supposing real possibilities than one who, in stricter spheres of reasoning supposes something self-contradictory and argues validity from it. Indeed we may, if we wish, think of each stretch of argument as preceded by a saving hypothetical clause, by such words as 'If such a being, or such a type of experience, were possible....'.[61]

60 With one important caveat that I will return to below.

61 Strawson, 1959. Pages 85-86.

I take this move to be definitive of Strawson's approach to sound perceptions. By abstracting from the role of the hearer's body, however, Strawson makes it difficult to know how to apply his account to any actual cases of sound perception. Sound perceptions become something “hypothetical” that are only possible “if such a being or such a type of experience were possible”. All actual sound perceptions, on the other hand, are had by hearers with a particular kind of body, and it would be strange if the body of the hearer had no effect on the nature of the sound perception.

In the next section, I discuss some of the ways in which the hearer's body bears on sound perceptions. As a result of the interaction of sound waves with certain features of the hearer's body, sound perceptions make the hearer aware of the location of a sound. Since Strawson omits the interaction of sound waves with the hearer's body, it is no surprise that he decides on a non-spatial view.

3.10: *Sound Perceptions and the Hearer's Body*

If it can be shown that the hearer's body plays a causal role in determining the nature of sound perceptions, then an understanding of that role is necessary for a complete account of sound perception. Strawson argues that sounds themselves do not contain the resources to produce in the hearer a concept of spatial location, but if the combination of sounds themselves and the body of the hearer *does* contain the resources for a concept of spatial location, there is reason to think that perceived location is in fact a feature of sound perceptions.

Thanks to much experimental research in hearing, the role of the hearer's body in determining the nature of sound perceptions is relatively well understood. The two most significant features of the body that affect the nature of sound perceptions are the position of the ears relative to the head and the rest of the hearer's body, and the shape of the ears.[62]

The position of the ears is significant in at least two respects. Firstly, because the ears are separated by a short distance, sound waves tend to reach one ear before the other. While sound waves move too fast for hearers to be conscious of this time difference, the difference is shown to produce in the hearer a conscious experience of direction. This can be seen in experiments where hearers wear a pair of headphones and receive an auditory stimulus in each ear separated by a time difference that would simulate that of sound waves reaching the ears in everyday hearing.[63] Remarkably, some listeners experience a difference of perceived direction given a time difference as small as one hundredth of a millisecond.[64]

Secondly, because the ears are not separated by an empty space, but by the head, a solid object which interferes with sound waves, the sound waves that arrive at one ear are not the same as the sound waves arriving at the other ear. If a sound is coming from the left, the sound waves that eventually arrive at the right ear will be of decreased intensity. For one, the sound waves must travel further to reach the right ear if the sound is coming form the left, but the head also dampens the intensity of sound waves as they come from the left, diffract around the head, and stimulate the right ear.[65] While the hearer is unaware of the difference in intensity at each ear, similar experiments to those measuring the

62 Culling and Akeroyd, 2010. Pages 123-130.
63 Ibid. Page 123.
64 Klump and Eady, 1956. Pages 859-860.
65 Brungart and Rabinowitz, 1999. Pages 1465-1479.

effects of time difference show that a difference in intensity produces in the hearer a conscious experience of sound direction.[66]

The shape of the ears also plays a role in generating sound perceptions. The majority of cues about where sounds are coming from that result form the shape of the ears are created by the way sound waves reflect within the concha, the “inner curl” of the ear.[67] The nature of the sound waves that go on to stimulate the inner ear depends on which part of the concha reflects those waves into the ear, which in turn depends on the initial direction of travel of the sound waves. Sound waves that travel from a higher elevation will be reflected by a lower region of the concha and *vise versa* such that sound waves are reflected from a unique region of the concha depending on their direction of travel. While hearers are obviously unaware of which part of their concha is reflecting sound waves, experiments show that the conscious experience of a sound's direction varies when the shape of the concha is distorted.[68]

The role that the body of the hearer plays in the generation of sound perceptions is important because it constantly underlies and shapes actual sound perceptions of creatures like us. By ignoring the way sound waves interact with the body, Strawson leaves out what seems to be the most important or even the sole origin of spatial information in sound perceptions. It is unclear whether Strawson leaves out the role of the body because he does not believe it to be relevant, or because he does not consider sounds to be sound waves. An important reason for suspecting the latter is that Strawson does not count the direction of travel or elastic nature of sound waves among the qualities of a sound. He

66 Hartmann and Constan, 2002. Pages 1037-1045.
67 Lopez-Poveda and Meddis, 1996. Pages 3248-3259.
68 Culling and Akeroyd, 2010. Page 127.

may have though such qualities to be inaudible, but the above experiments appear to suggest that such qualities are audible in the sense of producing the experience of perceived direction. In any cause, his omission of the way sound waves interact with the body of the hearer makes it is no surprise that Strawson concludes that sound perceptions do not contain any spatial information.

There is, however, one important caveat raised by both Strawson and Malpas, 1965. It seems that none of the information that is presented to the hearer as a result of the combination of sound waves and the hearer's body would count as *spatial* information if, in the first place, the hearer lacked a sense of the position of their own body.[69] Information provided by the combination of sounds and the hearer's body fixes the location of a sound relative to the hearer's body, but if the hearer has no sense of the location of their body to begin with, then sounds cannot be heard to have a location relative to the hearer's body. Sounds are heard to be “to the left” or “to the right”, but these descriptions of perceived direction are senseless until we know, “to the left *of what*” or “to the right *of what*”. Clearly what is meant by such a description is “to *my* left” or “to the left *of my body*”, but then how does information about the location of one's body – something which is obviously not a sound – get involved in sound perceptions? I will call this “the problem of egocentric location” (location assigned on the basis of a subject's location).

Malpas, 1965 states the problem of egocentric location in the following way:

69 Strawson, 1959. Page 66. Malpas, 1965. Page 141.

> . . . the location of sound necessarily makes reference to physical bodies. Now does this not suggest that there is something odd about saying that it is a feature of our auditory experience which enables us to locate sounds? We should have to say that this feature enabled us to detect a relationship between a sound and something which is not a sound, but a physical body, namely ourselves.[70]

Strawson is also motivated by this problem:

> Whatever it is about the sounds that makes us say such things as 'It sounds as if it comes from somewhere to the left', this would not alone (i.e. If there were no visual, kinaesthetic, tactual phenomena) suffice to generate spatial concepts.[71]

Whereas Malpas is quite explicitly dealing with the problem of how information about the location of the hearer's body is related to sound perceptions (and is doing so in reference to Strawson, 1959), Strawson states the problem as an argument in favor of the non-spatial view: the perceived direction of a sound is nothing without correlated “kinaesthetic” phenomena, which I take to be an awareness of the location or position of the hearer's body.

The nature of an awareness of one's body is outside the scope of this essay, but there are a few ways one might respond to Strawson and Malpas. Firstly, it might be said that a sense of the location of one's body is not a separate sense modality, like sight or touch, and it is not something external to hearing or sense perception in general, but rather a background feature of consciousness that is common to all conscious experiences including perceptual experiences.[72] If all perception takes place against a background awareness of the location of one's body, then it would be strange to think this awareness

70 Malpas, 1965. Page 141.
71 Strawson, 1959. Page 66.
72 Searle, 2004 seems to suggest such an approach (pages 98-99).

as something "external" to sound perceptions. One must also be awake to have a sound perception, but wakefulness is not considered to be external to perception; rather, wakefulness is usually taken to be a part of perceiving something.

One might also respond to the problem of egocentric location by pressing the claim that perceptions in general can make one aware of non-egocentric locations. Why not say that the perceived location of all of the objects of perception, whether sensed by hearing, sight, touch, etc., is egocentric? If all perceptions of location are egocentric, then there is no special problem of egocentric location for sound perception in particular.

To this one might object that vision presents objects as having an "allocentric location", a location relative not to the subject, but to other objects.[73] For example, a car might be seen to be located "in front of the house" or "beside the other car" rather than egocentrically, as "to my left" or "in front of me". This objection, however, masks the fact that everything in the visual field is perceived as having the location that it does as a result of the location of the subject's body, and in particular as a result of where the eyes are directed. Anything one sees can be *described* as having an egocentric location or an allocentric location, but visual perceptions themselves present the entire visual field relative to the location and position of the viewer's body, so any visible location is in the first instance egocentric or "relative to the subject's point of view".

On the other hand, if one were to insist that vision presents objects as having an allocentric location because, for instance, a car can be seen to be "beside another car", then why not insist that hearing presents sounds as having an allocentric location, as when the sound of one person's voice is heard to be beside the sound of another person's voice?

73 This point was suggested to me in a discussion with Professor O'Callaghan.

(Imagine that one is standing in a group of people talking; the fact that we do not normally speak of sounds as “beside” each other does not seem more persuasive than the experience that each person's voice seems to be coming from beside, above, or below every other person's voice).

Whatever one might conclude about the problem of egocentric location, it remains true that sound perceptions are dependent on more than just sounds themselves. The combination of sounds and the hearer's body provides more information to the hearer than could be derived from sounds alone. Theoretical considerations about sound perceptions that are based on the nature of sounds alone, like those of Strawson, 1959, are therefore insufficient to establish any conclusions about sound perceptions. Moreover, one might also complain that Strawson ignores the relevant features of sounds themselves, namely the ones responsible for how sound waves interact with the hearer's body, like direction of propagation or the elastic nature of sound waves.

3.11: *Practical Considerations for the Non-spatial View*

I think the main appeal of the non-spatial view is that it seems to have the potential to explain cases of hard to locate sounds. If sound perceptions do not present the hearer with any locational information, then it is no surprise that hearing is sometimes misleading or unhelpful about location. On the non-spatial view, this is because whenever the hearer is presented with a sound, the resulting sound perception must be correlated with outside information from the other senses or from memory. This opens the door to errors arising from correlated information or the process of correlation itself, and suggests that hearers will be unable to tell where sounds are coming from if no outside information is available.

3.12: *Hard to Locate Sounds*

Suppose a lecturer in a large, crowded hall hears somebody asking a question and also sees somebody stand up and start moving their lips. If the person asking the question is not the same as the person standing up and moving their lips, the lecturer will likely make a mistake about where the sound of the voice is coming from. The sight of somebody standing up and moving their lips gives misleading correlated visual information that leads the lecturer to make a mistake about where the sound is coming from. The non-spatial view might challenge spatial views by suggesting that, if sound perceptions present the hearer with a location, then the hearer should be able to tell where the sound is coming from quite independent of what they see.

According to the non-spatial view, cases where information from memory or the other senses is unavailable for correlation with sound perceptions also presents a case where hearers will be unable to tell where sounds are coming from. These cases are perhaps more common and more convincingly explained by the non-spatial view than cases where outside information is misleading. One group of cases where outside information about the location of a sound is unavailable is that of unfamiliar sounds. The hearer perceives a sound with a certain volume, pitch, and timbre quite clearly, but still cannot tell where the sound is coming from without seeing or otherwise perceiving the source.

For instance, if one walks into a room and hears a very strange high pitched sound, it could seem to one that the sound is coming from a squeaky radiator against the wall, from some squeaky pipes beneath the floor, from a strange bird outside the window, etc. In such cases, the location of the sound seems unclear. As O'Shaughnessy, 2009 suggests, "...make the sound an utterly unfamiliar noise. What is left of spatial perception?"[74] Due to the unfamiliarity of the sound and without any outside information from memory or the other senses, one cannot tell where the sound is coming from.

Pasnau, 1999 points to similar cases that can even arise when the sound is familiar. Some sounds, like that of a cricket somewhere in one's house, are notoriously hard to locate even though one perceives the pitch, timbre, and loudness quite clearly.[75] In tough cases like these as well as cases of unfamiliar sounds, the non-spatial view might claim, if sound perceptions represent the location of the sound, then the hearer should have no trouble telling where the sound is coming form.

Though cases where misleading outside information from memory or the other senses is correlated with sound perceptions and those where no such information is available can motivate the non-spatial view, they do not ultimately support the claim that sound perceptions contain no spatial information. This is easy to show in the first class of cases. Consider the example of the lecture hall where one person stands up and moves their lips, but another person is the one asking the question. This example, rather than suggesting that sound perceptions do no present any locational information, suggests that sound perceptions do not present location with a high degree of determinacy.

If the person asking the question is seated beside the person who stands up, then it will really seem that the person asking the question is in the same location as the person

74 O'Shaughnessy, 2009. Page 124.
75 Pasnau, 1999. Page 311.

standing up. Hearing alone does not present the fine grain of detail needed to differentiate the location of the person standing up from the person seated beside who is really doing the talking. But if the person who stands up is on the far left side of the hall and the person asking the question is on the far right, there will be little or no mystery about which side of the room the sound is coming from. Hearing alone will present the direction of the person asking the question as the far right side of the lecture hall. If hearing presented no spatial information whatsoever, this would not be so. While such cases show that hearing does not present locations in exact detail, it cannot be concluded from this that hearing presents no location at all.

Cases where information from memory or the other senses is unavailable for correlation with sound perceptions support the same conclusion. As O'Shaughnessy and Pasnau suggest, it can be very difficult to tell where an unfamiliar noise is coming from, or even where a familiar noise is coming from in some cases. But even O'Shaughnessy, who is skeptical about the presentation of locations in sound perception, does not make the strong claim that sound perceptions do not present locations at all: “Now make the sound an utterly unfamiliar noise. What is left of spatial perception? As it seems to me, direction and little else.”[76] Similarly, Pasnau observes that “Crickets are a notoriously difficult case, but even there you will hear the sound as having some general location. And it *is* possible to find a cricket in your house, just by listening to its sound. . .”[77] While such cases point to the indeterminacy of one's awareness of location, they do not support the conclusion that hearing presents no location at all.

76 O'Shaughnessy, 2009. Page 124.
77 Pasnau, 1999. Page 311.

3.13: *The Case Against the Non-Spatial View*

The above theoretical motivations for the non-spatial view (sections 3.9 and 3.10), while raising important questions about the nature of sounds and sound perceptions, ultimately hinge on larger issues than sound perception. Strawson's account of sound perceptions is the result of his view of sounds themselves and of perception in general. If Strawson were to count qualities like the direction of propagation or elastic nature of sound waves, he may well have concluded that sounds do have the qualities necessary to produce an idea of location. If his account of perception took consciousness with such background features as a sense of the location of one's body as components of perception, he may not have thought of perceived location as something that arises from “correlated” information.

If one accepts Strawson' basic approach, then much of what he says about sound perception would appear to be valid. In sections 3.9 and 3.10, I have tried to offer some reasons for rejecting Strawson's basic approach and for thinking his account of sound perception to be invalid, but there is much more to be said. Do sense modalities like hearing include background features of consciousness like wakefulness and an awareness of one's location as components, or are these external, correlated processes? If a sense of one's location is external to sense modalities, does this imply that all senses including even vision and touch are non-spatial (and should one be willing to accept such a conclusion)?

What is most appealing about the non-spatial view is its potential to explain cases where hearers cannot tell where sounds are coming from. However, since we cannot accept the non-spatial view's conclusion that, in such cases, sound perceptions present no locations at all, in rejecting the non-spatial view we are left wanting an alternative explanation of cases where hearers have difficulty telling where sounds are coming form. Whereas the remote view, in holding that sound perceptions make the hearer perceptually aware of the location of the source, attributes to the hearer more locational information than one actually has, the non-spatial view does not attribute enough.

It is important not to see the remote view and the non-spatial view as the only two alternatives. In discussions of sound perception, one sometimes gets the impression that if the remote view is false, sound perceptions must be devoid of any perceived location. Sorensen, 2007, for example, suggests that ". . . if anything is to be said about the location of a sound, it must be in terms of its source."[78] Consequently, it might seem that if one does not accept the remote view, one must hold a non-spatial view. In the rest of this essay, I present the medial view of sound perception as a middle way that can accommodate the insights of both views without inheriting their flaws. Though, against the remote view, sound perceptions do not make the hearer perceptually aware of the location of the source, sound perceptions do make the hearer perceptually aware of the location of the sound, a region of the medium which surrounds the hearer and the source.

78 Sorensen, 2007. Page 284.

IV

The Medial View

Inasmuch as sounds themselves are concerned, the most widely endorsed theory of sounds has been the view that sounds are sound waves in a medium.[79] By contrast, the medial view of sound perception enjoys little support. According to the medial view, hearers experience sounds to be located in the medium between the source and the hearer. If a bird is chirping outside the window, the sound of the bird is not heard to be fixed in place exactly where the bird is, but to more or less fill the space between the bird and the hearer. This chapter is dedicated to describing the details of perceived location according to the medial view and to showing how the medial view offers solutions to some of the problems that I have raised against the remote view and the non-spatial view.

4.1: *Perceived Location and the Boundary Problem*

In the last chapter, I argued that the model according to which perceived location is the conjunction of a perceived direction and perceived distance is problematic. The claim that sounds are heard to have a direction and a distance implies that we hear sounds to have boundaries. If a sound is heard to be ten feet away rather than nine feet away, the hearer must in some sense perceive the sound to have a boundary which begins ten feet away in

---

79 Casati and Dokic, 2010. sections 2.1-2.2. Pasnau, 1999 calls this part of the “standard view” of sounds (pages 309-310).

some particular direction.[80] I argued that this is a problem for the remote view because, on the remote view, the location of the sound is the exact same location as that of the source, which typically has *rigid* boundaries. Hearing, however, seems to be insensitive to rigid boundaries.[81] I cannot hear the sound of a car to occupy exactly the location that it does with anywhere near the level of determinacy with which I can see or feel the car to occupy that exact, rigidly defined location.

In this section, I argue that the medial view provides a model of perceived location which avoids the implication that sounds are heard to have rigid boundaries, or any boundaries at all. While a version of the medial view which holds that sounds are heard to have rigid boundaries could be formulated, I take it to be a virtue of the medial view that it can be formulated so as to avoid the boundary problem altogether. Because sounds are heard to fill the space stretching from the hearer's body to the sounding object, sounds are not heard to be “distant”; that is, sounds are not heard to be at any distance of remove from the hearer's body, like “nine feet away”, “ten feet away” etc.

As a spatial view, the medial view holds that sound perceptions have the form “*s* is at *p*”, but *p* does not include a perceived distance. The perceived location of a sound is a matter of perceived direction only, which raises two important questions: how can sound perceptions include a perceived direction without including perceived boundaries? And how can hearers learn the location of the source, especially the distance of the source, if perceived location is merely the perceived direction of the sound itself? (I address the

80 I emphasize that in this hypothetical, it is the sound that is heard to be ten feet away, not the source of the sound. The idea that a sound could be heard to be ten feet away is consistent with the remote view (or a version of the medial view which holds that we hear sounds to have robust boundaries), but the version of the medial view I defend does not claim that sound perceptions involve a perceived distance or any perceived boundaries at all.

81 Sorensen, 2007. Page 284. Pasnau, 2000. Page 38.

latter question in section 4.2).

The medial view's account of perceived direction avoids the boundary problem because the perceived direction of a sound is the direction in which a sound seems to be most greatly concentrated (to be loudest). Sounds seem to be concentrated in a certain direction relative to the hearer, but the hearer is not thereby made aware of the boundaries of a sound. If it sounds to me as if a car is approaching from my left, then I perceive the sound to lie to my left, but I do not perceive the size or the boundaries of the space that the sound occupies. After all, I might open my eyes, find that I am in a movie theatre, and realize that the apparent was actually the sound of some speakers fixed in place to my left. Whereas one might have *guessed* that the sound occupied a different region of space as the car seemed to approach, one did not perceive any such region of space since the actual space the sound occupied was a fixed space in the theatre. One perceived the direction of the sound, but not the dimensions or boundaries of the sound.

According to the medial view, the perceived direction of a sound is veridical if the greatest concentration of sound that is in earshot does in fact lie in the direction one experiences. On the medial view, the notion of the “concentration” of a sound figures into perceived direction because of the nature of the medium. If sounds are something in the medium, and most likely are sound waves or some properties of sound waves, then sounds, like sound waves, are something that occupy large regions of space. For example, if the sound of a speaker's voice were identified with the sound waves produced by the speaker, then the sound, like the sound waves, will be located throughout the entire room or auditorium, all around the hearer. In this sense, hearers are always “in the middle” of sounds.

Chapter IV

What is the perceived direction of a sound if sounds are something that surround the hearer? Although one is always in the middle of a sound, one does not usually experience sounds as coming from every direction (more on this phenomenon in section 5.6). Despite *some* concentration of sound in every direction around the hearer, sounds are not spread out evenly in every direction. Because sound waves dissipate as they travel away from the source, there is more sound nearer to the source and progressively less sound further away. Perceived direction, on the medial view, is our capacity to detect the direction in which sounds are most highly concentrated. Hearers sometimes perceive sounds to be all around them if the sound is very reverberant or very loud, but the perceived direction of a sound is almost exclusively directed toward the highest concentration of the sound.[82] Hence sounds are heard to be, for the most part, *between* the hearer and the source (or the hearer and the highest concentration of sound within earshot).[83]

Describing perceived direction in terms of the concentration of a sound avoids the boundary problem because hearing need not be sensitive to boundaries in order to be sensitive to the concentration of a sound. In principle, any view of sound perception could account for perceived location in terms of sound concentration. The remote view could be restated to say that the perceived location of a sound is the location of the greatest

82 In general, the auditory system is about 80-90% directed at the highest concentration of sound, according to Shinn-Cunningham, Zurek, and Durlach, 1993. The first sound waves to strike the ear, namely those which come directly from the highest concentration of sound (where the source is) are given precedence over sound waves that arrive at the ear indirectly from every other direction (as when sound waves reverberate against surfaces in the environment and are reflected into the ear). This is the "precedence effect", which is thought by auditory scientists to play an important role in perceived direction. (I infer that the auditory system gives precedence to the direction of the highest concentration of sound from the fact that it gives precedence to the direction of sound waves that travel in a straight line from the highest concentration of sound).

83 I limit the perceived direction of a sound to the highest concentration of that sound *that is within earshot* because obstacles can make sounds partly inaudible. The perceived direction of such a sound therefore depends on location of the greatest concentration of sound that the hearer is capable of perceiving.

concentration of sound, which, a proponent would have to claim, is the exact same location as that of the source. The flaw in such an approach, however, is that the source is located at a discreet distance, and one would have to perceive the distance of the concentration in order to perceive the location of the source. Such a view would again run into the boundary problem. On the medial view, however, adopting an account of perceived direction in terms of sound concentration allows proponents to give an account of perceived location which avoids attributing to hearers a perceived distance, and therefore avoids the boundary problem. This is only possible if sounds are heard to more or less fill the medium, and therefore not to be at any distance of remove from the hearer.

4.2: *Learning the Location of a Sounding Object*

Regardless of one's views about sounds or sound perceptions, sources (sounding objects) have discrete locations at a particular direction and distance from to the hearer. This seems to present the medial view with a significant challenge: if sound perceptions merely serve to make the hearer aware of the direction of a sound, how do we learn the direction and distance of the source, and how is it that we so readily learn where the source is just by hearing its sound? In rejecting the remote view, one rejects the most direct explanation of how hearers know where sounding objects are, as well as the explanation that seems to be most consistent with our everyday experience of hearing.

Normally when one hears a sound, one knows where its source is located immediately. Objections in the last chapter, however, suggest that, in spite of this knowledge, one does not *perceive* the location of the source. In the case of sounds in the presence of obstructions or the case of engineered sounds, one is easily fooled about the location of the source if hearing is all one has to go on. Whereas these cases create auditory illusions on the remote view, the medial view denies the claim that the location of the source is perceptible on the strength of hearing alone. Rather, hearers perceive qualities of sounds, some of which make the hearer perceptually aware of location (like perceived direction) and some which do not (like perceived pitch, loudness, timbre). All of these qualities give reliable cues to the location of the source. In perceiving these qualities, hearers are perceptually aware of the location of the sound and readily learn the location of the source in most cases, but hearers are not perceptually aware of the location of the source.

Two features of the medial view's account of how one learns the location of a sounding object stand out. Firstly, the medial view is emphatically not a description of the everyday experience of sounds. In section 1.3 I argued that this is a virtue. The everyday experience of hearing is the wrong place from which to draw an account of sound perception because, in everyday hearing, we are not interested in sounds themselves or sound perception. We are interested in and actively paying attention to sounding objects which we frequently also see or otherwise perceive. An account of perception must be wary of the subject's other mental operations, since it is notoriously difficult to differentiate perceptions from the influence of beliefs, desires, expectations, memories, attitudes, attention, and so forth. This is why the medial view can be motivated by less commonplace cases like that of hearing in the presence of obstructions or engineered sounds. The medial view should be thought of as a response to the question, “how does hearing alone function *in general*, i.e. in *all* of these cases?” rather than the question of how hearing functions in our everyday lives.

Secondly, by placing the hearer at one level of remove from perceptual awareness of the source's location (but not at any level of remove from the direction of the sound), the medial view accommodates error about the source's location without attributing auditory illusions to the hearer. It would be misleading to think of the remote view as the view of sound perception which explains how hearers successfully locate sounding objects, and the medial view as the view which explains how hearers fail. Rather, the most desirable view of sound perception will be the one which leaves room for hearers to make all of the familiar mistakes in locating sounding objects without attributing systematic illusions or defects to the sense of hearing. It could also be said in favor of the medial view that it is not so counterintuitive as to claim, with the non-spatial view, that sound perceptions do make the hearer perceptually aware of any location whatsoever. While the medial view puts the hearer at one level of remove from the location of the source, the non-spatial view creates a gulf.

4.3: *The Direction of the Source*

Learning the direction of a source is straightforward on the medial view. Since the medial view holds that hearers perceive the direction of the greatest concentration of sound, perceived direction typically coincides with the direction of the source: the greatest concentration of a sound is typically right in front of the source, so a straight line from the hearer to the greatest concentration of sound is usually also a straight line to the source. If a bird is chirping, the greatest concentration of sound is right in front of the bird's beak, so when I perceive the direction of the greatest concentration of sound, I perceive the direction of the bird.

While the remote view also holds that the perceived direction of the sound and the direction of the source coincide, this is usually conceived of in terms of the shared boundaries of the sound with its source. Since I hear the sound to be in the same rigidly bound location as the source, I know the direction of the source. For the medial view, however, the perceived direction of the sound coincides with the direction of the source because the greatest concentration of sound in the medium is (usually) in close proximity to the source.

In cases where there is an obstruction between the hearer and the source, the perceived direction of the sound is misleading about the direction of the source. When one is indoors, the sounds from outside seem to come from the direction of open windows quite independently of the direction of their source. Wherever a bird outside is chirping, the sound seems to come from the direction of an open window. If one accepts the remote view's claim that the perceived direction of a sound is the same as the direction of its source, there is no way to avoid attributing illusion to the hearer. The direction of the source is quite independent of obstacles in the environment, but if the perceived direction of the sound is equated with the direction of the source, then obstacles between the source and hearer will always create auditory illusions.

Such cases do not force proponents of the medial view to attribute auditory illusion to the hearer. The sound, at least insofar as it is audible, really is located in the direction of the open window (if the sound is taken to be something which more or less fills the medium). One will surely be misled about the location of the source, but one nevertheless accurately perceives the sound itself to be located in the direction of the window.

On the medial view, the effect of obstacles is to change the extent to which a sound is audible, thereby changing where the greatest concentration of audible sound is. While one and the same sound is indeed in greater concentration near the source, this does not create an illusion about the direction of the greatest concentration of sound if the part of the sound nearest the source is made inaudible. One simply perceives less of the sound. The perceived direction of a sound, even in the case of obstructions, is still the greatest concentration of audible sound. If birds are chirping outside, the greatest concentration of sound seems to be the open window, and this is in fact where the greatest concentration of sound that is audible to the hearer lies.

4.4: *The Distance of the Source*

Since perceived distance is not a feature of sound perceptions, the distance of a sounding object, unlike its direction, cannot be inferred from spatial qualities of sounds. Hearers learn the distance of a sounding object by perceiving certain qualities of its sound that provide reliable cues to the distance of the source. These can generally be classified as a sound's “acoustic qualities”. Perceiving a sound's acoustic qualities does not make the hearer perceptually aware of the distance of the source, but allows for it to be accurately inferred.

The best examples of acoustic qualities are loudness and reverberation.[84] The sound of someone shouting sounds much louder in a small, confined space than it does outdoors or

84 There is some ambiguity about loudness. On the one hand, we might speak of loudness as the output loudness of a sound at its source. Loudness as an acoustic quality, on the other hand, is the loudness of a sound *as it has been affected by its environment*. For example, loudness increases or decreases depending on whether materials in the environment reflect or absorb sound waves, and decreases as sound waves travel. The loudness of a sound as it is affected by its environment is the subject of this section.

in a much larger space. The perceived increase of loudness in a small space can be understood in terms of “ground effects”, the effects of soundwaves being reflected from relatively hard surfaces in the environment.[85] While some parts of sound waves travel directly from the source to the hearer, others travel from the source to the ground and are reflected from the ground to the hearer. Assuming that the ground is a hard, flat surface that will reflect rather than absorb sound waves, the effect on the hearer is a perceived increase in loudness of about three decibels as compared to the same sound in a space where there is no surface that reflects sound waves.

The sound of someone shouting from thirty feet away can also be heard to have different levels of reverberation depending on the space the sound occupies. The sound of someone shouting from thirty feet away in a large stone hallway or a subterranean cave has much more reverberation than the same sound outdoors. Perceived reverberation is like “echo”, but should not be confused with the special case where one shouts into an environment like the Grand Canyon, waits, and hears an “echo” of their voice after a period of silence. In contrast, reverberation is echo that is perceived simultaneously with the original sound, or almost simultaneously, since reverberating sound waves take longer to travel to one's ear than those which travel directly from the source.

If sounds are not heard to be fixed in place at a distance from the hearer, how is the hearer able to discover the distance of a sounding object? On the medial view, perceived acoustic qualities like loudness and reverberation function as cues that are systematically correlated with the distance of the source, allowing hearers to accurately infer the distance of the source. Since sound perceptions do not make the hearer perceptually aware of the

85 Davies, 2010. Pages 383-385

distance of the source, hearers must also depend on their knowledge of what acoustic qualities a sound *should* have in a certain space given that the source lies at any particular distance.[86] For example, if I know that I am in a fairly non-reverberant room, and I know what the sound of my dog barking is like in such a room when my dog is, say, ten feet away, then when my dog starts barking from this distance, I am easily able to infer that my dog is ten feet away on the strength of hearing alone.

How loud one hears a sound to be is the most important cue to the distance of the source. An intuitive way to describe the systematic correlation of loudness with the distance of the source is that a sound that is perceived to be loud typically has a source that is nearby, whereas the same sound with the source at a much greater distance is perceived to be relatively quiet. For example, if an explosion is two hundred feet away, the explosion sounds quieter than if it had been only twenty feet away, and louder than if it had been two thousand feet away. Perceived loudness therefore provides partial information about how far away a sound's source is, but is relatively unhelpful if the hearer is unfamiliar with the sound or its environment, and therefore does not know how loud it should seem if the source were at a given distance.

Reverberation also provides partial information about the distance of a sound's source. A sound that is perceived with a high amount of reverberation is likely to come from far away, whereas a sound perceived with little reverberation is likely to come from nearby (since the further sound waves travel, the more opportunity they have to reverberate against planes in the environment, if there are any). For example, the sound of someone

86 Culling and Akeroyd, 2010 goes into some detail about the role of the hearer's familiarity with the environment, the source, and the sound itself in learning the distance of a source (pages 128-130).

shouting at a hearer from the end of a long stone hallway is perceived to have a high amount of reverberation, whereas the same sound is heard to have very little reverberation if the person shouting is only a few feet down the hallway. As with perceived loudness, however, if the hearer is unfamiliar with the sound or with how much reverberation to expect from the environment, then perceived reverberation is relatively unhelpful in learning the distance of the source.

The reason perceived qualities like loudness and reverberation provide accurate information about the distance of a sound's source is that hearers are normally more or less familiar with the sounds they hear and the environments they hear them in, and therefore have accurate expectations about the distance of the source given a certain perceived loudness and reverberation. If one hears the sound of a person yelling and the sound seems quiet, hearers typically have the accurate expectation that the person yelling is relatively far away.

4.5: *Indeterminacy and Descriptions of Perceived Location*

The perceived location of a sound can be described relative to the hearer's body (egocentrically) or relative to the environment or objects in the environment (allocentrically). The sound of my television, for instance, can be described as “in front of me” or “all around me” (egocentric), or it can be described as “in the front of my living room” or “all around my living room” (allocentric). While both types of description pick out the same location, egocentric descriptions like “in front of me” describe the location of sounds with reference to the hearer's body. Allocentric descriptions like “all around my living room", on the other hand, describe sounds with reference to the environment or objects in the environment (and without reference to the hearer's body).

Because there is a degree of indeterminacy in hearing where sounds are located, one might think that sound perceptions do not present hearers with the level of detail required to attend to or describe sounds as having an allocentric location. According to this worry, sounds just seem to have a pitch, loudness, and timbre, and with the right combination of these qualities, we assign sounds a location *ex post facto*. If we were really able to perceive the location of a sound, we would have no trouble giving allocentric descriptions of a sound's location.

Malpas, 1965 gives this as a reason for thinking that sound perceptions might not make the hearer aware of any location whatsoever.[87] In section 3.10, I argued that sounds considered in isolation from the body of the perceiver need not cause a sound perception which makes the hearer aware of a location, but sounds in conjunction with the body of a normal human perceiver can in fact make the hearer aware of the location of a sound. In this section, I outline the sense in which sounds are heard to have an indeterminate location on the medial view, and how the indeterminacy of location figures into descriptions of perceived location.

The distinction between egocentric and allocentric that is so useful in descriptions of sounds does not map onto perceptual experience. Even though I can give an allocentric description of the sound coming from my television, like "the sound is in my living room", I cannot help but perceive the sound of my television from the location of my body (just like anything else I perceive). All perception, since it takes place from the

87 Malpas, 1965. Page 141-142.

location of the perceiver's body, is egocentric. By attending to the object of a perception in terms of its spatial relations to other objects and its environment, one is able to give an allocentric description of where things are located. But the things one perceives, along with their location relative to each other and the environment, is always perceived by way of the location of the perceiver's body. I see my television and hear its sound, and I may attend to the location of the television or that of its sound relative to the rest of the living room, but I am perceptually aware of all of this from the location of my body. In short, perceptual awareness is always egocentric, but descriptions of what one perceives can be egocentric or allocentric.

The worry that sound perceptions do not present hearers with the level of detail required to describe sounds as having an allocentric location can be addressed by looking at common allocentric descriptions of sounds. Beyond the problem of egocentric and allocentric location, the nature of allocentric descriptions of sounds also sheds some light on the indeterminacy of sound perceptions. Consider the following examples:

1. “The sound is coming from the basement!”
2. “That knocking sound is coming from the front/passenger side of the car.”
3. “That scraping sound is coming from the brakes.”
4. “Sound is coming from the right speaker, but the left one is broken.”
5. “The music is too loud in here!”
6. “The music is too quiet here, we should move closer to the stage.”
7. “The sound at the House of Blues is always great.”

All of these allocentric descriptions of sounds have one of the following forms (which is not to say that these are the only possible forms):

1. “The sound is coming from place *p*.”
2. “The sound is coming from object *o*.”
3. “It is loud/quiet in place *p*.”
4. “It is loud/quiet at place *p*.”

Chapter IV

In all of the above examples, a sound is described as having some location relative to an environment (“the basement”, “front/passenger side of the car”, “in here”, “at the House of Blues”) or objects in an environment (“the brakes”, “the right speaker”) rather than relative to a hearer's body.[88] On the strength of hearing alone, one could not give any of these descriptions because the place or object in each description is not something hearing alone could make one aware of. On the medial view, this is because hearing makes one aware of a particular direction, but all of the places or objects in the above descriptions are things that we identify by way of their boundaries. “The basement” and “the front/passenger side of the car” are meant to refer to locations with relatively well defined boundaries, but sound perceptions do not make the hearer aware of boundaries.

Nevertheless, sound perceptions do make the hearer aware of these locations, albeit in a way that they are not normally identified. If I hear a sound coming from the basement, I am not perceptually aware of the boundaries or dimensions of the basement (though I may be familiar with them from having seen them in the past). But I am perceptually aware of the direction of the basement and can therefore describe the sound as “coming from the basement”. Similarly, if I hear a sound coming from the front/passenger side of the car, I am not perceptually aware of the boundaries of the front/passenger side of the car. I could not sketch the boundaries just by hearing a sound coming from the front/passenger side. But, when I hear the sound, I am perceptually aware of the direction of the front/passenger side, and I can therefore give the allocentric description “the sound is coming from the front/passenger side”. Since I am not aware of the boundaries of the

88 An allocentric description of sounds relative to each other is also possible; e.g. “I hear that sound to be coming from below that other sound” (suppose a dog is barking at one's feet while a bird is chirping in a tree).

basement or of this region of the car, I am not perceptually aware of these locations as I would normally identify them (i.e. by seeing or otherwise perceiving their boundaries). I am, however, perceptually aware of these locations insofar as I am aware of their direction.

The peculiarity about allocentric descriptions of a sound's location is that locations in general are usually identified and talked about by way of vision, or in any case by way of their boundaries, but sound perceptions do not make one aware of boundaries. This points to the limitations of sound perceptions to make the hearer aware of locations. One is only made aware of the location's direction. Relative to vision, the capacity of sound perceptions to make one perceptually aware of locations is significantly diminished. But sound perceptions should not be held to the standards of vision at the expense of appreciating the extent to which sound perceptions do in fact make the hearer aware of locations. While sound perceptions do not make the hearer aware of the boundaries of a sound, they do make the hearer aware of a certain direction, and this is enough to locate sounds relative to their environment or objects in the environment.

4.6: *“Coming From” Descriptions and the Medial View*

Another peculiarity about common descriptions of sounds is the use of the phrase “coming from” in descriptions like “the sound is coming from the basement”. “Coming from” descriptions can be either allocentric descriptions, as in the previous sentence, or egocentric descriptions, like “the sound is coming from my left”. The difficulty with “coming from” descriptions is to understand the perceptual state(s) that hearers have when describing a perceived sound in this way.

O'Callaghan, 2009 distinguishes between a “spatial sense” and a “causal sense” of “coming from”.[89] In the causal sense of “coming from”, the causal origin of a sound is identified as such. For example, a speaker might use the expression “the sound is coming from the brakes” in order to identify the brakes as the causal origin of a sound. On the other hand, the spatial sense of “coming from” is used to predicate a location of a sound. Using the spatial sense, a speaker might say, “the sound is coming from the left”, in order to draw attention to the location of the sound. O'Callaghan, 2009 argues that because one does not hear sounds to travel, the causal sense of “coming from” is the only use of this expression that is consistent with sound perception:

> . . . sounds do not auditorily seem to travel toward us from their sources. . . Sounds therefore do not seem to come from their sources in any spatial sense of *coming from*. The sense in which it is correct to say that sounds seem to come from their sources must be a *causal* sense. Sounds seem produced or generated by their sources.[90]

On the medial view, the causal sense of “coming from” may seem problematic. If sounds are heard to be at one level of remove from their source, then what perceptual state does the hearer have when using a “coming from” description in the causal sense? The above definition of the causal sense of “coming from” is, however, consistent with the medial view: because most non-engineered sounds are a relatively reliable guide to the nature and location of their source, hearers can successfully communicate about sounding objects just by hearing their sound. If a sound seems to be coming from the brakes, it usually is coming from the brakes (but not always!). On the medial view, the

89 O'Callaghan, 2009. Pages 11-12.
90 Ibid. Cf. Scruton, 2009.

regularity of the nature of a sound with the nature and location of the source guarantees the utility of “coming from” descriptions in the causal sense.

However, the fact that hearers can communicate about sounding objects just by hearing their sound does not necessarily mean that sound perceptions make the hearer perceptually aware of sounding objects. If hearing a sound made the hearer perceptually aware of the nature and location of its source, hearers could not misidentify the source without being subject to an auditory illusion. Since the medial view puts the hearer at one level of remove from the source of a sound, misidentifying the nature or location of a source is understood as a mistaken inference rather than an auditory illusion.

On the medial view, “coming from” descriptions can also be used meaningfully in the spatial sense, and in a way that does not imply that sounds are heard to travel. According to the medial view, sounds are heard to fill a region of space with indeterminate boundaries between the hearer and the greatest concentration of sound (which is usually in front of the source). In the spatial sense of “coming from”, where a location is predicated of a sound, the speaker is perceptually aware of the direction of the greatest concentration of sound and wishes to communicate this direction without any care for the source or its location. The spatial sense of “coming from” is less common than the causal sense because, in day to day life, we are typically if not exclusively interested in sounding objects rather than their sounds. If I say, “the sound is coming from the brakes” or “the sound is coming from my left”, the sound is usually the last thing I am interested in. I want to know what is wrong with the brakes or what might be lurking to my left.

Nevertheless, if the location of the sound is one's concern, then “coming from” descriptions in the spatial sense can be used in a way that is consistent with the medial view to communicate the perceived direction of the sound. But how does one get from the perceived direction of a sound to the idea that the sound is actually *coming from* that direction? On the medial view, hearers do not get from a perceived direction to the idea that sound are coming from a location on the strength of hearing alone. One might know or believe that a sound is coming from a location, but one cannot perceive that a sound is coming from a location (at least not on the strength of hearing alone).

The idea that sounds come from a location comes from the experience that sounds become quieter the further away one gets from the greatest concentration of sound.[91] This experience, however, requires much more than hearing. One only hears the loudness of a sound at the location where one is standing, so learning that sounds becomes quieter as one moves away from their point of greatest concentration requires one to consciously walk around, and to identify the sound one hears as the same sound in a weaker form as one walks away. From this complex experience, one learns that sounds grow weaker or "dissipate" outward from their point of greatest concentration. When one says "the sound is coming from my left" in the spatial sense, one means that the greatest concentration is located to one's left and that the sound is therefore dissipating outward from one's left – but on the strength of hearing alone, one is only perceptually aware of the direction of the greatest concentration of sound and loudness of the sound at one's particular location (among other auditory qualities like pitch and timbre).

91 As Perkins, 1983 notes, the fact that sounds are "perceptually weakened" by ones distance from the source suggests that sounds might in some sense travel (on his view), but this is *inferred* from the weakening of the sound and does not imply that hearers *perceive* sounds to travel (page 164).

Chapter IV

4.7: *The Role of the Hearer's Body*

It is an interesting feature of the medial view that features of the body discussed in the previous chapter, like the position of the ears on the head and the shape of the ears, have a "direct" explanation: on the medial view, sounds are heard to be located in the medium, and therefore to be *something* in the medium.[92] This makes the role of the relevant bodily features "direct" in the sense that the body of the hearer is in direct physical contact with the medium, and therefore in direct physical contact with sounds. For example, if sounds are sound waves or systems of sound waves, then there is no causal intermediary between sounds and the hearer's body.

The modern scientific account of perceived location runs parallel to the medial view. On the scientific explanation, the perceived location of a sound is explained by the interaction of the hearer's body with the medium, and with sound waves in particular. The position and shape of the ears, which are thought to be the most important or even the sole bodily features in fixing a perceived location, are only relevant to hearing where things are given the interaction of sound waves in the medium with the hearer's body. For example, the majority of cues about where sounds are coming from that result form the shape of the ears are created by the way sound waves reflect within the concha (the "inner curl" of the ear).[93] Likewise, on the medial view, since sounds are heard to be something in the medium, perceived location is also explained by the body of the hearer in

92 If sounds are heard to be located in the medium, it follows that sounds are heard to be "something" in the medium. I go into some detail about what sounds themselves might be assuming that the medial view of sound perception is correct in sections 5.1-5.5.

93 Lopez-Poveda and Meddis, 1996. Pages 3248-3259.

conjunction with features of the medium. Like the modern scientific explanation, the medial view explains the role of bodily features in perceived location without reference to anything beyond the way in which the medium interacts with the hearer's body.

The remote view does not share this parallel with the modern scientific view of perceived location. On the remote view, perceived location has an “indirect” explanation: if bodily features like the shape of the ears work in conjunction with the medium to generate a perceived location, yet sounds themselves are something located at the source, then sounds are not in direct physical contact with the body of the hearer to generate a perceived location. On the remote view, the medium “gets in the way” of perceiving sounds: sounds are located at the source and affect the medium, which then acts as a causal intermediary, carrying information about sounds more or less reliably to the hearer. Whenever this information is distorted, as in the case of hearing in the presence of obstructions, the remote view attributes auditory illusion. On the medial view, however, sounds are heard to be features of the medium, so it can be said that sounds directly affect the hearer's body, and that sounds are therefore directly responsible for their perceived location.

4.8: *Hard to Locate Sources and Unfamiliar Sounds*

In section 3.12, I argued that hard to locate sources and unfamiliar sounds (including sounds heard in unfamiliar environments) might motivate skepticism about the level of detail with which sound perceptions make the hearer aware of locations, but such cases do not support the claim that sound perceptions do not make the hearer perceptually aware of any location at all. Since unfamiliar sounds are generally problematic because they make the source hard to locate, I discuss cases of unfamiliar sounds and cases of hard to locate sources as one kind of challenge for views of sound perception. In the following section, I intend the expression "hard to locate sources" to encompass both of these cases.

Hard to locate sources can be shown to make a good case in support of the medial view. Hard to locate sources like the tough case of the cricket are instances where most will agree that *some* locational information is available. O'Shaughnessy, 2009 suggests that in such cases, the hearer perceives ". . . direction and little else".[94] Pasnau, 1999 suggests that, "you will hear the sound as having some general location. And it *is* possible to find a cricket in your house, just by listening to its sound. . .".[95] From this we might conclude that, rather than making the hearer aware of no location at all, sound perceptions make the hearer aware of a relatively indeterminate location. In tough cases like that of the cricket, one perceives some degree of directionality and perhaps little else. A cricket's sound may not seem to be coming from exactly 45° the left, but by paying attention to the sound's direction, the hearer can eventually find its source.

How does the remote view handle hard to locate sources? Tough cases like that of the cricket seem to be precisely the cases where a view of sound perception must hold that sound perceptions present relatively indeterminate locations. The location of a sounding object, however, is always a determinate location with fairly rigid boundaries. Consequently, the remote view appears to be unable to handle hard to locate sources.

94 O'Shaughnessy, 2009. Page 124.
95 Pasnau, 1999. Page 311.

A proponent of the remote view, however, might wish to accommodate hard to locate sources by suggesting that sound perceptions make the hearer aware of the location of the source, but in a relatively indeterminate way. On this modification of the remote view, sounds are perceived to be “in the neighbourhood” of their source. In the rest of this section, I outline three reasons for thinking that this modification is unavailable to proponents of the remote view.

The first reason is that the remote view is put forward as a theory about how hearers successfully locate the source of a sound and navigate their environment. If sound perceptions make the hearer aware of locations which are indeterminate with respect to boundaries, as the medial view holds, then the remote view seems to have no advantage over the medial view as a view about how hearers navigate their environment. On the medial view, hearers perceive sounds to be located in a region of space between the hearer and the source that is indeterminate with respect to boundaries, and use this and other auditory qualities to infer the exact location of the source.[96] If sounds or their environments are unfamiliar to the hearer, then auditory qualities that would otherwise have been helpful, like loudness and reverberation, will not help the hearer to locate the source. On the remote view, however, if it is stipulated that the hearer is perceptually aware of the location of the source as an indeterminate location with respect to boundaries, how does the hearer learn the exact location of the source? If the exact location of the source is not perceived, it must be inferred from other auditory qualities, but inferring the exact location of the source from the apparently indeterminate location

96 On the medial view, auditory qualities like loudness, reverberation, and clarity also figure into hearers ability to infer the location of a sounding object.

of the sound places this modification of the remote view into the same position as the medial view when it comes to navigating one's environment, namely the position of having to infer the location of the source.

Secondly, the suggestion that sounds are presented to be at the source in a relatively indeterminate manner seems to pose a phenomenological challenge. Suppose we have an upward facing speaker. On this modification of the remote view, the sound of the speaker would be perceived to be located in an indeterminate region of space all around the speaker.[97] The problem with this suggestion is that sounds have a greater intensity nearer to the source, and more intense sounds render less intense sounds inaudible (a phenomenon referred to by auditory scientists as “masking”).[98] Consequently, it is prima facie unlikely that hearers can perceive sounds to be all around their source. The sound as it is directly above the upward facing speaker will mask the sound on the side of the speaker furthest from the hearer because the sound directly above the speaker is more intense than the sound surrounding the far side of the speaker.[99]

Thirdly, the suggestion that, on the remote view, sounds could be heard to lie in a relatively indeterminate region of space around the source begins to collapse the remote view into the medial view.[100] If, on the remote view, sounds are not heard to be exactly at the source, then they are heard to lie in some region of the medium, perhaps in addition to lying at the source itself. While the medial view does not hold that sounds are heard to lie

97 This seems to be the view expressed in O'Callaghan, 2010. Page 124-125.

98 Oxenham and Wojtczak, 2010. Pages 7-10.

99 Alternatively, one might suggest that sounds are not more intense near the source but only seem to be so, and instead have a uniform intensity such that no masking occurs. I take this suggestion to be even less desirable as it makes the intensity of sounds, insofar as it varies by the distance of the source, a pervasive auditory illusion. Pasnau, 1999 seems content to accept this form of illusion (see pages 312-313). I return to this point in section 2.4.

100 Casati and Dokic, 2010 makes a similar observation (section 3.2.3).

in a region of space overlapping the source, it does take the perceived location of sounds to be indeterminate with respect to boundaries, so this distinction between the remote view and the medial view would be very weak should this modification of the remote view be adopted. Moreover, if the modified remote view is to avoid attributing auditory illusion to hearers, sounds would also have to be individuated as something which lies in the medium around the source and at the source itself. In the face of theses challenges, it would seem that the remote view must all but collapse into the medial view in order to account for hard to locate sources.

Hard to locate sources are best explained by the medial view because these are precisely the cases where the location of a sound is perceived to have indeterminate boundaries, and the medial view appears to be the only view that can accommodate this indeterminacy for systematic, non-*ex post facto* reasons: sources can be hard to locate because their sounds, like all sounds, are perceived to have indeterminate boundaries. If a sound is also unfamiliar, this only exacerbates the problem since, on the medial view, auditory qualities are used to infer the location of the source on the basis of the hearer's past experience with the way a sound should seem in a certain environment if its source is at a given location. This attempt to modify the remote view, on the other hand, runs into a least the three above problems: the motivation to give a view that explains hearers' ability to precisely locate sounding objects is defeated; the notion that a sound can seem to be in the space around its source is phenomenologically dubious since the greater intensity of the sound at the source will mask the sound on the far side of the source; and the remote view collapses into the medial view if it is construed as claiming that sounds are heard to occupy regions of space in the medium that merely happen to overlap with the location of the source.

4.9: *The "Indirectness" of Hearing a Sounding Object*

One of the deepest motivations for the medial view comes from a phenomenon noted by philosophers of all of the views so far discussed, and which might be referred to as the "indirectness" of hearing sounding a object. Perkins, 1983 describes this phenomenon in the following way:

> . . . an odour, like a sound, often comes to us without our seeing or otherwise knowing for certain what its source is. . .[101]

Pasnau, 2000 gives a similar account:

> . . . our view about sound seems to stem primarily from the fact that we see colors only when we look directly at them, whereas we hear sounds around the corner, down the hall, etc. This makes it seem as if sound fills the air, hence exists in the air, whereas color seems located in a single place.[102]

In both of these definitions, sounds are described as something that is often perceived without also perceiving the sounding object. In these cases, like when one hears a sound that is coming from around the corner, one is perceptually aware of the sound, but not of the source. As Perkins, 1983 says, we perceive the sound "without our seeing or otherwise knowing for certain what its source is". The sound of a conversation taking place around the corner of a building could be produced by people talking to each other, or it could be produced by a speaker or a set of speakers. On the strength of hearing alone, the hearer is not perceptually aware of what the source is. Pasnau, 2000 remarks, I think

101 Perkins, 1983. Page 163.
102 Pasnau, 2000. Page 38.

correctly, that these are obvious cases where sounds seem to "fill the air" (the medium) and do not seem to be located "in a single place" (but to be located in an indeterminate region of the medium).

I call this phenomenon the indirectness of hearing *sounding objects* because, according to the medial view, there is nothing indirect about hearing *sounds* in such cases. Sounds really do fill the air and are in direct contact with the body of the hearer.[103] Sounding objects, on the other hand, are at some distance of remove, both literally and epistemically. Sounding objects are not in direct contact with the body of the hearer, nor does one know anything for certain about a sounding object just by hearing its sound.[104]

According to Pasnau, 2000, the indirectness of hearing sounding objects is not a desideratum for theories of sound perception. Despite the fact that one sees colours only when one sees coloured objects, but hears sounds without hearing or otherwise perceiving sounding objects, sounds ought to be treated like colours: ". . . it is hard to see why this should be a satisfactory basis for saying that the one exists in the object, the other in the air. It rather seems merely a feature of the different physical properties of light versus sound waves."[105] Since sound waves bend around their environment and remain intact but light must be reflected directly into the eye, the phenomenology of hearing is different from that of vision; but why should this matter to a theory of sound perception?

Pasnau, 2000 is surely right to explain the phenomenological difference between hearing and vision as an effect of the different behaviour of sound waves and light, but

103 Since Pasnau, 2009 holds a remote view of sound perception and individuates sounds as events located at or near the source, the hearer is also said to be at some distance of remove from sounds themselves.

104 With sounds, on the other hand, one is perceptually aware of (and in this sense knows) the sound's audible qualities, its pitch, loudness, timbre, and so forth. Naturally, this does not mean that one can describe the audible qualities one hears, which would require some skill at describing sounds.

105 Pasnau, 2000. Page 38.

this does not trivialize the resulting phenomenological difference. If the fact that sound perceptions put the hearer at a level of remove from sounding objects is not reflected in one's theory of sound perception, then one's theory attributes auditory illusions: cases of hearing in the presence of obstructions and cases of engineered sounds show that sound perceptions are not sufficient to make the hearer perceptually aware of the location of the source, and cases of engineered sounds also show that sound perceptions are not sufficient to make the hearer aware of what the source is; if a theory of sound perception omits the indirectness of hearing sounding objects, then these become cases of auditory illusion.

I think there are two intuitive reasons that the importance of the indirectness of hearing sounding objects gets downplayed. I suspect that making these reasons explicit is enough to deflate their appeal. The first is that in day to day life, we are not interested in sounds, but in sounding objects.[106] We use sounds to get information about sounding objects and to navigate our environment. Even at an orchestra, one pays some attention to watching the musicians play, and so is perceptually aware of the location of the players by way of vision. In such conditions, even at an orchestra, one is liable to mistake the location of the sound for the location of its source. This slip is all to easy to make when one is perceptually aware of the location of a sounding object by way of vision: if this is the location that one is paying attention to, it is no surprise that one intuitively assigns this location to the sound as well. And this mistake becomes even easier to make given that sounds seem to have unbounded locations, which are therefore difficult to pin down in the first place.

106 This is discussed in more detail in sections 1.3 and 4.2.

The second reason, which is closely related to the first, is that the purpose of hearing in the evolutionary long haul, and even in our day to day lives, is to give information about sounding objects and the environment. If this is what the sense of hearing is adapted to do, then it seems hearing must make us perceptually aware of sounding objects. This account of hearing, however, ignores the role of the other senses. There is no justification for thinking that hearing alone must shoulder the burden of making hearers aware of sounding objects rather than merely playing a part in concert with the other senses. If sounds can give relatively reliable information about sounding objects, then hearing can play an advantageous evolutionary role without making the hearer perceptually aware of sounding objects. It can be seen from the fact that one would have no idea what a sounding object is if hearing were one's only sense that hearing is insufficient to make the hearer perceptually aware of sounding objects.

Both of these motivations against the indirectness of hearing where sounding objects are seem to be present in Pasnau, 1999, to take one example.[107]

> How do we manage to hear an orchestra, if not by hearing the sound of its performance? (Could our sensations of sound really be just epiphenomenal experiences, floating above the serious work of listening to objects in the environment?)[108]

Pasnau, 1999 demonstrates our day to day interest in sounding objects (as opposed so sounds themselves) by asking how we get information about orchestras rather than their sound. But one might respond, what do orchestras have to do with sounds anyway? An indistinguishable sound could be produced by speakers, and on the strength of hearing

107 Other examples are Kulvicki, 2008 (e.g. page 9) and O'Callaghan, 2010 (e.g. pages 123-124).
108 Pasnau, 1999. Page 317.

alone, there would be no perceived difference. One is not perceptually aware of a group of people holding musical instruments and moving their arms and legs in concert just by way of one's ears.

Pasnau, 1999 also calls our natural interest in sounding objects the “serious work” of listening, which suggests that the purpose of hearing is to make hearers aware of sounding objects. Gathering information about sounding objects is certainly serious work, but there is no justification for thinking that hearing alone must shoulder the burden. Sounds are a relatively reliable guide to their source and there is much value to be found in the regularities that exist between sounds and their sources, like the correlation between the loudness of a sound and the distance of the source.

In section 2.3, I also discussed Pasnau, 1999's claim that the duration of sounds is not a desideratum for a theory of sound perception. Not only do we often perceive sounds in the absence of their source, sounds are also perceived for different durations than all or most of the source's qualities. The colour, shape, size, taste, smell, etc. of objects are typically perceived to be constant over long periods of time and do not appear to change until the object is modified or damaged. A speaker maintains its colour, size, and shape until I modify or damage it. Sounds, on the other hand, are not perceived *until* an object is modified, and are typically perceived to endure for only seconds or minutes. The speaker always seems to have the same colour and shape, but it only sounds when it is modified by an electrical current causing vibrations. Similarly, other sounding objects do not make a sound until they are made to vibrate (sometimes by collisions with other objects).

Pasnau, 1999, argues that the fact that sounds endure for shorter periods of time than other qualities like shape or colour is not relevant to locating or individuating sounds.[109] But if sounds are not identified with something that endures for short periods of time, but rather with properties of objects that are constant for most of the object's existence, then one is left with the peculiar view that one is always surrounded by sounds in the same way that one is always surrounded by shapes and colours, yet one cannot hear these sounds except on rare occasion. If one is standing in a completely silent library, one would forced to say that one is surrounded by all sorts of sounds, none of which one can hear. Only be speaking or dropping one's books on the floor would one be able to uncover the sounds that are concealed all around one. On the medial view, however, sounds can be identified with properties of the medium, which is unstable and changing. By placing sounds at one level of remove from material objects and their properties, the medial view can individuate sounds in a way that is consistent with the intuition that sounds are not stable, constant features of the world.

I take it to be a virtue of the medial view that it avoids downplaying the indirectness of hearing sounding objects; avoids attributing auditory illusion to the hearer where the remote view does attribute illusion; and avoids the non-spatial view's skepticism about the capacity of hearing to make one aware of locations in any way at all. There are, however, important objections to the medial view. I now turn to these.

109 Pasnau, 1999. Page 322.

V

Objections to the Medial View

The most widely discussed objections to the medial view have a common theme. If sounds are heard to be located in the medium, then sounds must be identified with something that is located in the medium in order to avoid attributing widespread illusion about perceived location. Sound waves are the most obvious candidate, but sound waves behave in ways that are inconsistent with the experience of a sound. And if sounds cannot be identified with sound waves, what is there for proponents of the medial view to identify sounds with?

I respond to three objections in this chapter, each of which points to a way in which the behaviour of sound waves appears to be inconsistent with the experience of a sound, and which conclude that the medial view must therefore attribute widespread auditory illusion. This is a serious challenge to the medial view because abandoning sound waves altogether would seem to leave nothing for proponents of the medial view to identify as the (material) object of a sound perception.[110]

Three ways in which the behaviour of sound waves seems to be inconsistent with the

110 Even the suggestion that sounds might be identified with dispositional properties of the medium appears to be inconsistent with auditory experience. If sounds are identified with disposition properties of the medium, then sounds are properties that the medium always has. This is problematic because, as Pasnau, 1999 (page 322) points out, we experience sounds to be fleeting relative to the objects of vision or touch. It would be contrary to our experience of sounds to hold that sounds are constant features of the medium even though nobody can hear them (until an object interacts with the medium in the right way). A dispositional analysis of colours, on the other hand, is consistent with our experience. Objects always seem to have one colour or another, and even in the dark we insist that objects retain their colour.

experience of a sound are with respect to the motion, direction, and duration of sound waves. If one accepts the wave theory as it is commonly understood, one attributes auditory illusion with respect to at least these three features of sounds. Sounds do not seem to travel toward the hearer like sound waves do, nor to have the direction of the entire body of sound waves, which in fact surrounds the hearer, nor to have the duration for which the body of sound waves exist.

Rather than abandoning the wave theory in the face of these inconsistencies, I propose a revised version. Sounds are not identical with individual sound waves, but with properties of sound waves that emerge when groups sound waves exhibit certain patterns over time.[111] I call this the "holistic version" of the wave theory. In this chapter, I describe the problems posed by the motion, direction, and duration of sound waves, and elaborate the holistic wave theory in response. If the holistic wave theory is correct, then the motion, direction, and duration of sound waves is not an obstacle to identifying sounds with something in the medium. Consequently, the medial view cannot be refuted on the ground that there is nothing in the medium that it would be consistent with our experience of sounds to identify as the material object of sound perceptions.

5.1: *The Motion Problem*

Sound waves travel. If sounds are nothing but sound waves, then the fact that sound waves travel through the medium might be thought to have important implications for the medial view. One might think, for instance, that if sound waves travel, then on the medial view of sound perception, sounds ought to be perceived to travel. If not, then the

111 Nudds, 2009 gives a similar account (pages 75-77).

medial view attributes to the hearer the illusion that sounds are stationary when in fact they are as mobile as sound waves are. O'Callaghan, 2010 illustrates this point: if sounds are sound waves, then sounds, like sound waves, should be heard to approach the hearer, fill the hearer's head, and whizz past like “auditory missiles”.[112] The medial view, however, does not claim that we hear sounds to travel. Sounds, on the medial view, are heard to “fill spaces”, not to travel toward the hearer like an “auditory missile”. But how can sounds be heard to fill a space if, in fact, they are travelling through that space? This could be called the “motion problem”.

O'Callaghan, 2007's point that sounds are not heard to travel despite the movement of sound waves suggests that the wave theory of sound, as it is commonly conceived of, stands in need of revision. If what we hear are indeed individual sound waves travelling toward the hearer, then sounds should be heard to have the motion of the individual sound waves that one hears. I think the necessary revisions to the wave theory can be made by looking into two models of sounds in the history of philosophy that are still widely used in contemporary science and sound engineering, namely vibration and pressure. While much has been said about vibration and pressure with respect to sounds, I choose some observations made by Berkeley and Aristotle respectively as examples because they capture certain intuitions about sounds that I wish to highlight. I also take it to be a virtue of choosing accounts from the history of philosophy that such accounts are unbiased by a common conception of the wave theory today, which I take to be the view that sounds are identical to (individual) sound waves.

112 O'Callaghan, 2007. Pages 35-36.

5.2: *Vibration, the "New Philosophy", and the Motion Problem*

Descartes, characterizing the standard philosophical view of sounds, says that "most philosophers maintain that sound is nothing but a certain vibration of air which strikes our ears."[113] This raises many questions: if sound itself is vibration, how are sounds to be individuated? Which vibrations or which sound waves count as an individual sound? Is the first wave moving through the air the sound? The first five or ten waves? Whichever waves are striking the ear?

The problem with all of these suggestions is that the sound waves named in each case travel from the source to the hearer. Yet the view that vibrations or sound waves are constantly on the move *despite* the perception of sounds as stationary is, historically, the most widely held view of sounds in the philosophical literature. Since at least Aristotle, sounds themselves have commonly been described as something moving, but, to my knowledge, nobody has ever held the view that we experience sounds to travel from the source to the hearer.

Berkeley, in describing the position of the "new philosophy" associated with Boyle, Newton, and Locke, says that,

> A bell struck in the exhausted receiver of an air-pump [i.e. a vacuum] sends forth no sound. The air, therefore, must be thought of as the subject of sound . . . When any motion is raised in the air, we perceive a sound greater or lesser, according to the air's motion; but without some motion in the air, we never hear any sound at all. . [sound] is merely a vibrative or undulatory motion [in] the air.[114]

113 Descartes. *The World*. Page 82.
114 Berkeley, George. 1713. First Dialogue.

Chapter V

Berkeley's description of sounds according to the "new philosophy" emphasizes the role of air movement in bringing about sound perceptions. While the object of a sound perception is a body of moving air, sound are not perceived to have the motion of the corresponding body of air. Rather, sounds are perceived to have certain auditory qualities *in proportion* to the movement of bodies of air: ". . .we perceive a sound *greater or lesser, according to the air's motion*" (my emphasis). The motion, or more specifically the *vibration*, of a body of air is also said to be a necessary condition for a sound perception; ". . .without some motion in the air, we never hear any sound at all. . . [sound] is merely a vibrative or undulatory motion [in] the air".[115]

In conceiving of the motion of bodies of air as vibration, this account of sounds might seem to avoid the motion problem altogether. Bodies of air are said to "move" in the sense of vibration rather than to travel from the source to the hearer like an auditory missile. If sounds do not travel from the source to the hearer, then it is no surprise that in sound perception, sounds do not seem to travel.

Ultimately, however, the view described by Berkeley does not avoid the motion problem. If sounds are identical to vibrating bodies of air, then the fact that we do not experience sounds to vibrate implies that our sound perceptions are illusory in this respect. If sounds are said to travel from the source to the hearer *or* to vibrate, the hearer is subject to an illusion about the motion of sounds (since vibration is a kind of motion). But while this conception of sounds is not a solution to the motion problem, one of its features is worth highlighting.

115 Vibration may even be a sufficient condition for sounds in Berkeley's view. I say necessary because, presumably, vibrations of a medium that are outside the audible spectrum would not count as sounds.

The auditory qualities we perceive, like pitch, loudness, timbre, and reverberation, are said to *vary* by the amount or "proportion" of vibration in the medium; vibration is not itself a perceived auditory quality. The "new philosophy" did not have the motion problem in mind, but this is perhaps reason enough to wonder whether sounds were meant to be identified with vibrating bodies of air, or if sounds were meant to be identified with proportions of vibration in the medium. In any case, if sounds are identified with the amount of vibration, which we might understand as the frequency and/or amplitude of vibrations in the medium, then there may be a way to solve the motion problem. I return to this idea after considering a related idea in Aristotle's theory of sounds.

5.3: *Pressure, Aristotle, and the Motion Problem*

Whereas both the "new philosophy" and Aristotle connect sounds to movements of the medium, Aristotle also understands sounds in terms of the pressure of the medium.

> What has the power of producing sound is what has the power of setting in movement a single mass of air which is continuous from the impinging body up to the organ of hearing. . . Air in itself is, owing to its friability, quite soundless; only when its dissipation is prevented is its movement sound.[116]

The pressure created when the medium is prevented from dissipating is, according to Aristotle, a necessary condition for sounds.[117] Moreover, Aristotle calls sounds ". . . a single mass of air which is continuous. . .", which I take to mean that sounds, being the movement of bodies of air, are for Aristotle unified by having a continuous distribution of

116 Aristotle. *De Anima*. Book II, chapter 8, 420a0-420a11.

117 Pressure may even be a sufficient condition for sounds in Aristotle's view. I say necessary because, presumably, levels of pressure that are outside the audible spectrum would not count as sounds.

pressure.[118]

Today, we would state Aristotle's point about pressure a bit differently, but I think that Aristotle's definition of sounds is compatible with that of modern science. The standard definition of sounds in auditory science is that sounds are to be identified with sound waves.[119] For auditory scientists, this means that sounds are distributions of pressure in a medium. Davies, 2010, for example, defines sound waves as "... regions of increased and decreased pressure" in a medium.[120] Increased *and* decreased pressure because sound waves are understood to be fluctuations of pressure produced by vibrations rather than merely increased *or* decreased levels of pressure. Aristotle's remarks about pressure, the "new philosophy's" observations about vibration, and the conception of sound waves in modern science can be connected by examining how vibration relates to pressure.

When an object vibrates, it moves back and forth very fast. When it moves forward, it packs air molecules together, creating higher air pressure, and when it moves backward, it pulls air molecules apart, creating lower air pressure. Due to the elastic nature of a medium like air, the result is a body of air that is unified in the sense of having a distribution of pressure that is continuous (but not uniform) throughout; that is, a body of air whose pressure at any given point is a function of the distribution of pressure throughout.[121] The connection between vibration and

pressure is that the vibration of sounding objects brings about a vibration or "oscillation" of air molecules which, at the level of air molecules, can be described as a vibration, or at

118 Johnstone, forthcoming provides an alternative interpretation of Aristotle, connecting Aristotle's conception of sounds much more closely to sounding objects than to the medium.

119 Casati and Dokic, 2010. Sections 2.1-2.2. Pasnau, 1999 calls this part of the "standard view" of sounds (pages 309-310).

120 Davies, 2010. Page 375.

the level of whole bodies of air, can be described as a distribution of pressure. Whereas vibration is predicated of molecules, pressure is predicated at the higher level of whole bodies of air.

Observations from the "new philosophy" and from Aristotle supply rich conceptual resources for understanding sounds. Vibration and pressure, as anticipated in these traditions, are importantly related to sounds. While modern science goes into intricate detail about the many patterns that sound waves tend to exhibit, vibration or the oscillation of pressure in a medium is fundamental to our understanding of what a sound wave is. In the next section, I argue that sounds can be identified with vibrations or fluctuations of pressure in a medium if these are understood as higher level properties that belong to whole systems of sound waves.

5.4: *Higher Level Properties of Sound waves and the Holistic Wave Theory*

Recall that the problem with a common understanding of the wave theory of sound is that any individual sound waves travel from the source to the hearer in a way that the hearer is perceptually unaware of. Sound waves travel like auditory missiles, but this is not reflected in the experience of a sound. The insight that can be found in revisiting fundamental conceptions of vibration and pressure is that what one is aware of when one hears a sound might be said to be properties of whole systems of sound waves over a

121 As opposed to a body of air where the pressure at one point has no effect on the pressure at some other point. An example would be sounds in two soundproof rooms; in this case, rather than having a body of air with a continuous distribution of pressure, one has two separate bodies of air that each have their own continuous distribution of pressure, one within each room. Consequently, on the contemporary scientific view, one has two separate sounds rather than one continuous sound. Here one might accept that there are two separate sounds, but reject the claim that sounds can be identified with bodies of pressurized air or with individual sound waves.

period of time, like their frequency of vibration or fluctuation of pressure, rather than qualities of individual sound waves, like their motion, or the position of their crests and troughs (their points of increased and decreased pressure).

On the holistic view, individual sound waves merely make the hearer aware of systems of sound waves. Individual sound waves strike the ear, setting off a chain reaction that produces in the hearer a sound perception, but individual sound waves are not the object of which the hearer is aware. When one hears a sound, one is aware of a certain pitch, loudness, timbre, and so forth. Everyone with a normally functioning sense of hearing will be familiar with the experience of such qualities, even if one cannot differentiate pitch from other qualities, and even if the technical definition of these qualities alludes one. Whatever feature of the external world answers to this experience, and is therefore the material object of sound perceptions, it must be something which has the experienced qualities of pitch, loudness, and timbre.

The feature of the external world that can be said to have a pitch, loudness, and timbre can be seen by a brief examination of how these qualities are conceived of in contemporary auditory science. De Cheveigné, 2010, a survey of the current scientific literature on pitch, defines pitch as ". . . the rate at which a periodic waveform repeats itself".[122] In other words, sound waves get bunched up into collections which have a uniform pattern relative to each other, and when these collections are related in the right way (when they repeat themselves), the system in which these collections are found can be said to have a pitch. Individual collections of sound waves do not have a pitch any more than individual sound waves within a collection have a pitch; pitch is a property that

122 De Cheveigné, 2010. Page 71.

the larger system of sound waves inherits from the arrangement of its constituent sound waves. Similarly, certain arrangements of molecules will reflect light, causing the perception of a colour. But the molecules are not themselves coloured. Colour is a property that larger objects inherit from the arrangement of their constituent molecules.

In modern auditory science, “loudness is most often displayed or described as a function of physical sound intensity or pressure”.[123] Pressure, itself, behaves as a unified system. The behaviour of a body of air pressure which gives rise to a sound perception, for example, is determined by the vibration of a sounding object, the size, shape, and composition of the acoustic environment, and the ambient pressure and temperature of the medium. Changing any of these factors changes the behaviour of the entire system, which in turn determines the level of pressure at any given point within the system.

The particular level of pressure that determines perceived loudness is the level of pressure at the location of the hearer. The closer the hearer is to the source, the greater will be the pressure of the system of sound waves, and the greater will be the perceived loudness of the sound. However, the particular level of pressure at any given region of the system of sound waves does not exist independently of the distribution of pressure throughout the system of sound waves. Rather, the pressure at any point is a function of the distribution of pressure throughout the system.

Timbre is a property of systems of sound waves in a manner similar to that of pitch. The patterned collections of sound waves that constitutes a sound's pitch tend to occur at different frequency multiples of each other, and the relative loudness of these frequency

123 Epstein and Marozeau, 2010.

multiples is a sound's timbre.[124] Since we are again talking about relations of collections of sound waves, timbre is a property of the system to which these collections belong, rather than a property of any individual sound waves taken separately. Since pitch, loudness, and timbre are what the hearer is aware of in hearing a sound, it can be said that hearers are aware of properties belonging to systems of sound waves, and therefore that whole systems of sound waves are the material object of a sound perception.[125] This is the core of the holistic wave theory that I propose as an answer to the motion problem.

5.5: *The Motion Problem and the Individuation Problem*

How does the holistic wave theory solve the motion problem? Motion is a property of individual sound waves, rather than systems of sound waves. The first sound wave of a larger system travels from the source to the hearer, as does the second, and so forth; but the system of sound waves does not thereby travel. The region of space occupied by the system is by definition the total region of space occupied by *any* of its parts. As they develop over time, systems of sound waves grow and shrink in the same way that people or trees grow in height over time. But systems of sound waves do not move outward from the source any more than people or trees, in maturing over time, “move” toward the sky. Rather, the boundaries of the system change as it unfolds over time.[126]

Since systems of sound waves occupy regions of space rather than moving through the

124 Davies, 2010. Pages 376-377.

125 Nudds, 2009 suggests that a theory of this type comes closest to accommodating the ways in which we normally individuate sounds (pages 75-77).

126 Dretske, 1967. With respect to motion, systems of sound waves are like events. Events, unlike individual objects, are not wholly present at any one time (by definition). Consequently, in the case of an event, there is no one subject which, at time *t*1, is at place *p*1, and at time *t*2, is at place *p*2. The first part of the event takes place at *p*1, and the second part at *p*2, but neither part constitute the entire event, which would have to be wholly present at one time or another in order to do any travelling.

air like individual sound waves do (like auditory missiles), the holistic version of the wave theory avoids the motion problem. The sound of a bird chirping, for example, seems to roughly coincide with the location of the system of sound waves the bird creates, not to travel toward the hearer and whizz past like individual sound waves do. In formulating the holistic wave theory, I have also offered an answer to the individuation problem: in the causal chain leading up to a sound perceptions, sounds can be individuated as higher level properties which belong to whole systems of sound waves.

5.6: *The Direction Problem*

When something makes a sound, the product is a system of sound waves which fills a large region of space all around the hearer. If one is having a conversation indoors, individual sound waves produced by each speaker travel outward in every possible direction and are reflected off of all of the surrounding walls.[127] But this seems to be inconsistent with the experience of a sound. Sounds seem to have well defined directions that we typically describe as “to the left” or “to the right”, and that we could represent by pointing or by drawing a straight arrow pointing in the perceived direction of the sound. I will refer to this apparent inconsistency as “the direction problem”.

The direction problem is closely related to the boundary problem discussed in chapter IV. On the medial view, sounds are not heard to have boundaries. Moreover, the holistic wave theory does not attribute rigid boundaries to systems of sound waves. Surely systems of sound waves begin and end somewhere, but systems of sound waves dissipate as their distance from the source increases, making their boundaries “loose” and

127 Davies, 2010. Pages 379-382.

indeterminate. Does the system of sound waves end when, walking away from the source, the sound becomes inaudible? Or does the system end when the ambient pressure of the medium is no longer disturbed by the motion of individual sound waves, even if by that point the sound has become inaudible?

Whether one is inclined to say that systems of sound waves have indeterminate boundaries or no boundaries at all, it is not a highly controversial claim that sound perceptions do not make the hearer aware of the boundaries of a system of sound waves. Whereas Sorensen, 2007 finds this problematic to the claim that sounds can be heard to have a location independent of that of the source, proponents of the medial view may offer the following response (which I have outlined in section 4.1): the perceived location of a sound is a matter of perceived direction, and perceived direction is a perceptual awareness of the direction of the greatest concentration of audible sound. This suggestion accommodates the perceived direction of a sound while avoiding the claim that sounds are heard to have boundaries. Hearers do not perceive the location of a sound by way of its boundaries, but by way of their sensitivity to the concentration of a sound. The perceived direction of a sound is quite definite because the point at which a sound is most greatly concentrated is perceptually available to the hearer.

What about sounds that seem to come from every direction? And if systems of sound waves are more or less concentrated all around the hearer, why don't sounds always seem to come from every direction? The first question points to a somewhat uncommon case, like that of some loud rock concerts or a symphonies. On the medial view, if a sound seems to come from every direction, it would follow that the point of greatest concentration is spread out all around the hearer, and in such cases this is what one finds.

Chapter V

At a loud rock concert, the acoustics of the hall are typically engineered to reflect sound waves such that, between the sound waves that travel directly to the hearer and the sound waves that are reflected off of the surfaces of the hall, there is a more or less uniform distribution of sound waves all around the hearer. This creates the effect that the sound of the performance is not a weak, thin sound coming from the stage, but a full, rich sound that fills the hall. The sound of the concert seems to come from every direction because the greatest concentration of sounds is spread out in every direction around the hearer.

The second question is directed at the common experience of sounds. If there is always some level of concentration all around the hearer, then why does the perceived direction of a sound always seem to be the specific direction of the greatest concentration of sound? This phenomenon is explained by the weakness of the concentration of sound which actually does surround the hearer, as well as the role of attention and the role of the hearer's body.

The concentration of sound that surrounds the hearer, and especially the concentration of sound on the far side of the hearer (relative to the source), is significantly diminished. As sound waves travel further from the source, reach the hearer, and surround the hearer, their intensity is significantly reduced. If perceived direction is a function of the concentration of sound, then the perceived direction of the concentration of sound on the far side of the hearer is likewise diminished. Because the concentration of sound which actually does surround the hearer is much weaker than the greatest concentration of sound, the perceived direction of the weaker concentration of sound is also overpowered by that of the greater concentration. Interestingly, however, the perceived direction of the weaker concentration of sound on the far side of the hearer seems to be *vastly* overpowered by that of the greater concentration. Perceived direction is not proportional to concentration, but seems to be almost exclusively directed at the greatest concentration. I think there are two reasons for this.

Firstly, on the medial view, sounds are heard to fill the air, and part of hearing sounds to fill the air is hearing sounds to lie all around one. Since, however, one is typically interested in the source of a sound, one instinctively listens for the direction of the greatest concentration, usually to the exclusion of the more or less diminished concentration of sound that lies all around one. In everyday hearing, the direction of the lesser concentration of sound that surrounds the hearer is usually not recognized because hearers are actively attending to the direction of the greatest concentration (which is usually taken to be the direction of the source).

Secondly and more importantly, as Shinn-Cunningham, Zurek, and Durlach, 1993 have shown, the auditory system is about 80%-90% directed at the sound waves which strike the ear first.[128] This is known as the “precedence effect” or the “law of the first wavefront”. It is a feature of hearers' bodies that, when it comes to perceived direction, the auditory system largely ignores the sound waves that surround the hearer and instead processes the sound waves that travel directly from the point of greatest concentration (which is usually right in front of the source).

For these two reasons, perceived direction is not proportional to the concentration of sound, and is almost entirely a matter of the greatest concentration of audible sound. Normally, hearers are exclusively interested in the direction of the greatest concentration (since this is usually the direction of the source), and, even if one were interested in the

128 Shinn-Cunningham, Zurek, and Durlach, 1993. Pages 2923-2932.

direction of lesser concentrations of sound, the auditory system barely processes the direction of the sound waves that do not come from the point of greatest concentration.

5.7: *The Duration Problem*

O'Callaghan, 2007 states what I call the duration problem in the following way:

> Clearly, perceiving the durations of sounds is an important part of auditory perception. Sounds inform us about happenings in and states of our environment, and part of what they inform us about is how long those happenings and states last. I learn through hearing when the coin stops spinning, when the fridge starts up and shuts down, and how long the car idles in the driveway. . . Now, if sounds are spatially bounded particulars whose locations in the medium change from moment to moment as do those of waves, what in fact I experience when I take the sound to have duration is not the duration of a sound at all. Rather, my encounter with a *spatial* boundary of a sound leads to my enjoying an auditory experience while the sound passes. On later encountering the far boundary of the sound, I experience the sound to end. . . I mistake the duration of an experience alone for the duration of a thing I am experiencing. Duration perception is a wholesale illusion if sounds are waves.[129]

When the boundaries of a body of sound waves begin to strike the ear, one begins to hear a sound, and when the body of sound waves ceases to strike the ear, one ceases to hear the sound. The problem is that we take the period of time while we are experiencing a sound to be the duration of the sound, but if the wave theory is correct, this is merely the duration for which sound waves are striking the ear. Sounds themselves continue to exist for as long as the sound waves do. Rather than learning the duration of the sound, one merely learns the duration for which one perceives the sound. Consequently, O'Callaghan, 2007 concludes that the wave theory gives rise to pervasive auditory illusions about the duration of sounds.

129 O'Callaghan, 2007. Page43-44.

One could deny, however, that durations are something that can be perceived in the first place. 'Duration' does not name a perceptible quality like shape, colour, sound – nor, perhaps, any *quality* at all. One cannot point to a thing's duration, nor does one have any sense organs that are adapted to making one aware of durations. On his list of perceptible qualities in book II.6 of *De Anima*, Aristotle does not include duration or time as a proper sensible, common sensible, or incidental object of perception. In Book II.14 of *An Essay Concerning Human Understanding*, Locke states that the idea of duration arises from reflection rather than perception.[130]

Rejecting the claim that durations can be perceived eliminates the possibility of illusory duration perceptions that might be thought to arise on the wave theory, but this leaves us with two important questions. Firstly, as O'Callaghan, 2007 suggests, learning the duration of events in one's surroundings is an important function of hearing. One often learns how long the phone was ringing or how long the car outside was idling by way of hearing its sound. What epistemological role does hearing play in such cases if not to make the hearer perceptually aware of durations? How else could hearing sounds give the hearer accurate information about the duration of events in one's surroundings? Secondly, we sometimes speak as if to attribute durations to sounds: “The sound of the ring tone lasts for three seconds, then repeats”; “The song is three minutes and one second long”. If we do not perceive durations, then how should attributions of durations to sounds be interpreted?

Hearing accurately informs the hearer about the duration of events in one's surroundings because hearing enables one to *infer*, but not perceive their duration. The

130 Locke, 1690. Book II,14.

duration of a sound perception is always the same as the duration for which a sounding object sounds, despite the fact that the sound waves will continue to endure after travelling past the hearer. For as long as an object is sounding, sound waves strike the hearer's ear; if an object sounds for five seconds, the hearer will have a sound perception that lasts for five seconds. This regularity between the duration soundings and the duration of sound perceptions allows hearers to reliably infer the duration for which the object sounds. But sounding is something the source does, rather than a feature of its sound. The epistemological role of hearing is not to make the hearer perceptually aware of durations (since durations are not qualities and cannot be the material object of a perception), but to provide the hearer with a reliable foundation to infer the duration of events in one's surroundings.

The fact that sound perceptions inform the hearer about the duration of events in their surroundings is reflected in ascriptions of duration to sounds. As Perkins, 1983 suggests, most ascriptions of duration to sounds are actually intended to ascribe duration to events in one's environment (to soundings rather than to sounds):

> Of course for most occasions on which we have an interest in noticing how long a sound lasts we are chiefly interested in how long the sound-making object persists in its sound-making action. (How long did the telephone ring? How long did the dog continue to whine?) And about this our auditory perception informs us correctly.[131]

When we talk about the duration of a sound, we are usually interested in the duration for which an object sounds rather than the duration for which the sound exists (the duration for which the systems of sound waves exist). Perkins' view points to an

131 Perkins, 1983. Page 172.

ambiguity in the way we talk about sounds: usually by "duration of the sound", we mean "duration of the sounding" or "duration for which the object sounds". This is certainly consistent with the fact that in everyday hearing, we are almost exclusively concerned with sounding objects rather than sounds themselves. The questions, "how long was *the sound* of the telephone ringing?" or "how long was *the sound* of the dog's whining?" would sound a little strange in ordinary language (as if there might be some deception about how long the phone itself was ringing).

For cases of everyday hearing, this sounds right. If one asks, "how long did the telephone ring?", one typically wants to know something about the phone rather than its sound; namely, how long was the phone making the sound or how long was the person calling trying to get through. If one asks, "how long was the dog whining for?" or "how long did you hear the dog whine for?", one wants to know how long the dog was making the sound for, or how long the person perceived the sound and had to put up with it. One is either asking about the duration for which the object sounds or the duration for which the hearer perceived the sound, but not the duration of the sound itself.

But what about cases outside of everyday hearing, or cases where one is simply not asking about the duration for which an object sounds? O'Callaghan, 2007 gives the following example:

> It is simply a mistake according to the wave account to state that 'Time is on My Side' by the Rolling Stones is three minutes and one second long if the song is the sounds.[132]

The song "Time is on My Side" really is three minutes and one second long

132 O'Callaghan, 2007. Pages 44-45.

(depending on the version). But if the song is the sound waves, then it would seem that the song is actually whatever period of time the sound waves endure for. It is clearly false that the song is whatever period of time the sound waves endure for, so it would appear that the wave theory cannot accommodate the duration of sounds themselves.

Rather than rejecting the wave theory, one might instead reject the clause which the above passage depends on, namely ". . . if the song is the sounds". This clause seems to rest on a type/token ambiguity. If "the song" is taken to be a token sound, which on the wave theory is a token collection of sound waves, then the wave theory would have the absurd consequence that the duration of the song is otherwise than we know it to be.

But songs are not token sounds. Performances of songs produce token sounds, but songs themselves are sound-types. This can be seen from the fact that songs can be written and never performed. A song may have a title, a chord progression, and a duration, but never once be instantiated in a performance. And if sound-types like songs are completely independent of token sounds, then the duration of sound-types is completely independent of the duration of token sounds.

Whereas the behaviour of sound waves fixes the duration of token sounds, the duration of sound-types is simply a matter of definition. "Time is on My Side" is three minutes and one second long because the Rolling Stones decide that it is. How this decision comes about, however, deserves some attention.

A sound-type may get a duration by simply declaring it to be so. If I write the song, I get to decide how long it is. But sound-types may also be defined by way of performances. If my band writes a new song on the spot, then we may decide that the duration of the song will from now on be the duration of that particular performance of the song, and in this case the sound-type gets the duration of the event which produced it; that is, the sound-type is given the duration for which the source sounds (and not the duration of a token collection of sound waves). In addition to performances, sound-types may also be defined by way of sound perceptions. If someone hears my band playing a five minute long song and says, “the first three minutes were great, but it should end there”, then we might decide to make the duration of the song three minutes. In this case, the sound-type gets the duration of the listener's sound perception (and again not the duration of a token collection of sound waves).

Since ascriptions of duration to sounds are either ascriptions of duration to sound-types, to events of objects making sounds, or to sound perceptions – and not to token sounds (collections of sound waves, on the wave theory) – the wave theory does not imply any problematic conclusions about ascriptions of duration. In everyday hearing, one is interested in sounding objects and what they are doing, so it is no surprise that most ascriptions of duration to sounds are actually ascriptions of duration to soundings. Other ascriptions of duration are to sound-types, or to sound perceptions (the *intentional* object of which we usually just call a “sound”, making ascriptions of duration to sound perceptions easy to mistake for ascriptions of duration to token sounds). And since durations are not perceived, but arise from reflection, the wave theory cannot imply any auditory illusions about duration.

5.8: *Conclusion*

While the motion problem is solved by conceiving of sounds as higher level properties of systems of sound waves, the direction and duration problems are solved by elaborating and clarifying the medial view. In the case of the direction problem, sounds are perceived to have definite directions because hearers usually attend to the point of greatest concentration, since this is usually indicative of the direction of the source. Lesser concentrations of sound which fill the air around the hearer are also overpowered by the greater concentration of sound, and are more or less ignored by the auditory system as a matter of brute fact (and likely as a consequence of the evolutionary advantage of directing attention to the greatest concentration of sound, which is usually located in the direction of the source). Nevertheless, in circumstances like that of a rock concert in a good concert hall, the concentration of sound is more evenly distributed and sounds really do seem to come from every direction and fill the air.

In the case of the duration problem, there is no illusion about duration because duration are not perceptible. The experience of a sound has a duration which the hearer is certainly conscious of, and the duration of that experience aligns with the durations one might be interested in, from the duration of soundings to the duration of sound-types. If a song is three minutes and one second, then one experiences a performance of that song (a sounding) for just that duration, and therefore also has an experience with the same duration as the song itself (a sound-type).

Of the three views of sound perception considered in this essay, the medial view offers the most solutions and suffers the fewest objections. This is not to say that the medial view is without problems or, of course, that it has been fully explored. The problem of how to count sounds or of how to distinguish sounds from each other, for example, has not been solved. I merely hope to have shown the medial view to be the best of the available alternatives insofar as they have been considered.

Bibliography

Anscombe, G. E M. “The Intentionality of Sensation”. *Analytical Philosophy*. 2nd ed. Ed. R. J. Butler. Oxford: Blackwell, 1965.

Aristotle. *De Anima*. *Introduction to Aristotle*. Ed. R. McKeon. Trans. J. A. Smith. New York: Random House, 1974.

Austin, J. L. *Sense and Sensibilia*. Ed. G. J. Warnock. New York: OUP, 1962.

Berkeley, G. *Three Dialogues Between Hylas and Philonous in Opposition to Sceptics and Atheists.*1713. Accessed July 30, 2013 <http://www.gutenberg.org/files/4724/4724-h/4724-h.htm#second>.

Brungart, D. S. and Rabinowitz, W. M. “Auditory Localization of Nearby Sources: Head -related Transfer Functions”. *Journal of the Acoustical Society of America*. Vol. 106. 1999.

Casati, R. and Dokic, J. “Sounds”. *Stanford Encyclopedia of Philosophy*. Ed. Edward N. Zalta. 2010. Accessed Jul. 01, 2013 <http://plato.stanford.edu/archives/win2012/entries/sounds/>.

———— “Some Varieties of Spatial Hearing”. *Sounds and Perception: New Philosophical Essays*. Ed. M. Nudds and C. O'Callaghan. Oxford: OUP, 2009.

———— “Sounds”. *Stanford Encyclopedia of Philosophy*. Ed. Edward N. Zalta. 2005. Accessed Jul. 01, 2013 <http://plato.stanford.edu/archives/win2005/entries/sounds/>.

Culling, J. F. and Akeroyd, M. A. “Spatial Hearing”. *The Oxford Handbook of Auditory Science*. Ed. Christopher J. Plack. Vol. 3. New York: OUP, 2010.

Davies, W. J. “The Acoustic Environment”. *The Oxford Handbook of Auditory Science*. Ed. Christopher J. Plack. Vol. 3. New York: OUP, 2010.

De Cheveigné, A. “Pitch Perception”. *The Oxford Handbook of Auditory Science*. Ed. Christopher J. Plack. Vol. 3. New York: OUP, 2010.

Descartes, R. *The World*. *The Philosophical Writings of Descartes*. Vol 1. Ed. Trans. J. Cottingham, R. Stoothoff, and D. Murdoch. Cambridge: CUP, 1985.

Dretske, F. E. “Can Events Move?” *Mind*. Vol.76(304). 1967.

Bibliography

Epstein, M. and Marozeau, M. “Loudness and Intensity Coding”. *The Oxford Handbook of Auditory Science*. Ed. Christopher J. Plack. Vol. 3. New York: OUP, 2010.

Evans, G. “Things Without the Mind”. *Philosophical Subjects: Essays Presented to P. F. Strawson*. Ed Z. Van Straaten. Oxford: Clarendon, 1980.

Hartmann, W. M. and Wittenberg, A. “On the Externalization of Sound Images”. *Journal of the Acoustical Society of America*. Vol. 99. 1996.

Hartmann, W. M. and Constan, Z. A. “Interaural Level Differences and the Level-meter Model”. *Journal of the Acoustical Society of America*. Vol. 112. 2002.

Hatfield, G. “Perception as Unconscious Inference”. *Perception and the Physical World: Psychological and Philosophical Issues in Perception*. Ed. D. Heyer and R. Mausfeld. Chichester: John Wiley & Sons, 2002.

Helmholtz, H. L. *Handbuch der physiologischen Optik*. 1867. 3rd ed. Ed. A. Gullstrand, J. von Kries, and W. Nagel. Leipzig: L. Voss, 1910.

Johnstone, M. “Aristotle on Sounds”. *The British Journal for the History of Philosophy*. Vol.21(5). 2013.

Klump, R. G. and Eady, H. R. “Some Measurements of Interaural Time Difference Thresholds”. *Journal of the Acoustical Society of America*. 28. 1956.

Kulvicki, J. “The Nature of Noise”. *Philosopher's Imprint*. Vol. 8(11). 2008.

Locke, J. *An Essay Concerning Human Understanding*. 1690. Ed. John Yolton. London: Aldine House, 1961. Vol. 1.

Lopez-Poveda, E. A. and Meddis, R. “A Physical Model of Sound Diffraction and Reflections in the Human Concha. *Journal of the Acoustical Society of America*. Vol. 100. 1996.

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

Malpas, R.M.P. “The Location of a Sound”. *Analytical Philosophy*. 2nd ed. Ed. R. J. Butler. Oxford: Blackwell, 1965.

Moore, D. R. “Series Preface”. *The Oxford Handbook of Auditory Science*. Ed. Christopher J. Plack. Vol. 3. New York: OUP, 2010.

Nave, C. R. *HyperPhysics*. Accessed Jul. 01, 2013 <http://hyperphysics.phy-astr.gsu.edu/hbase/waves/standw.html#c2>.

Bibliography

Nudds, M. “Sounds and Space”. *Sounds and Perception: New Philosophical Essays*. Ed. M. Nudds and C. O'Callaghan. New York: OUP, 2009.

O'Callaghan, C. “Perceiving the Locations of Sounds”. *Review of Philosophy and Psychology*. Vol. 1(1). 2010.

———— “Constructing a Theory of Sounds”. *Oxford Studies in Metaphysics*. Vol. 5. 2009.

———— *Sounds: A Philosophical Theory*. New York: OUP, 2007.

O'Callaghan, C and Nudds, M. “Introduction: The Philosophy of Sounds and Auditory Perception”. *Sounds and Perception: New Philosophical Essays*. Ed. M. Nudds and C. O'Callaghan. New York: OUP, 2009.

O'Shaughnessy, B. “The Location of a Perceived Sound”. *Sounds and Perception: New Philosophical Essays*. Ed. M. Nudds and C. O'Callaghan. Oxford: OUP, 2009.

———— *Consciousness and the World*. Oxford: Clarendon, 2000.

———— “Seeing the Light”. *Proceedings of the Aristotelian Society*. Vol. 85. 1984.

———— “Material Objects and Perceptual Standpoint”. *Proceedings of the Aristotelian Society*. Vol. 65. 1964

———— “The Location of Sound”. *Mind*. Vol. 66(264). 1957.

Oxenham, A. J. and Wojtczak, M. *The Oxford Handbook of Auditory Science*. Ed. Christopher J. Plack. Vol. 3. New York: OUP, 2010.

Pasnau, R. “The Event of Color”. *Philosophical Studies*. Vol. 142. 2009.

———— “Sensible Qualities: The Case of Sound”. *Journal of the History of Philosophy*. Vol. 38(1). 2000.

———— “What is a Sound?” *The Philosophical Quarterly*. Vol. 49(196). 1999.

Perkins, M. *Sensing the World*. Indianapolis: Hackett, 1983.

Reid, T. *An Inquiry Into the Human Mind*. 1764. Ed. Jonathan Bennett. 2007. Accessed Jul. 01, 2013 <http://www.earlymoderntexts.com/rein.html>.

Russell, B. *The Problems of Philosophy*. 1912. Oxford: OUP, 2001.

Bibliography

Scruton, R. "Sounds as Secondary Objects and Pure Events". *Sounds and Perception: New Philosophical Essays*. Ed. M. Nudds and C. O'Callaghan. Oxford: OUP, 2009.

Searle, J. *Mind: A Brief Introduction*. New York: OUP, 2004.

Shinn-Cunningham, B. G., Zurek, P. M., and Durlach, J. I. "Adjustment and Discrimination Measurements of the Precedence Effect. *Journal of the Acoustical Society of America*. Vol. 93. 1993.

Sorensen, R. *Seeing Dark Things: The Philosophy of Shadows*. New York: OUP, 2007.

Strawson, P. F. *Individuals: An Essay in Descriptive Metaphysics*. 1959. New York: Routledge, 1990.