

Against Hearing Meanings

Casey O'Callaghan

Philosophical Quarterly, forthcoming

Do you hear the meanings of utterances spoken in a language you know? At first blush, there are two ways to interpret the question. The first takes 'hear' loosely. It concerns whether upon hearing an utterance you become aware of its meaning or semantic properties. This way of interpreting the question takes 'hear' loosely because, for example, hearing an utterance, understanding it, and grasping its meaning in thought entails an affirmative answer. The first interpretation is not my focus. The second interpretation, which is the focus of this paper, does not take 'hear' so loosely. It concerns whether, in hearing an utterance spoken in a language you know, you are auditorily perceptually aware of its meaning or semantic properties. This is the interpretation I intend when I ask: Do you hear meanings?

Some philosophers have suggested the answer is 'yes' on the grounds that some episodes of auditory perceptual experience involve awareness as of meanings or semantic properties of spoken utterances. I am skeptical. In this paper, I explore and ultimately reject one motivation for accepting that humans hear meanings. The motivation, in short, is that awareness of meanings has a distinctive impact on the phenomenal character of auditory perceptual experience. The argument for this is that hearing meanings is part of the best explanation for the phenomenal contrast between listening with understanding to speech in a language you know and listening to speech in an unfamiliar language. If so, meanings are audible features that affect the phenomenal character of auditory experience—they are among the apparent features with which perceiving subjects are auditorily acquainted. I argue that the phenomenal character of auditory perceptual experience when listening to spoken language does not justify belief in hearing meanings. If meanings or semantic properties are among the contents of auditory perceptual experience, they are not among those that constitutively shape its phenomenal character. Meanings thus do not figure in

any content of auditory perceptual experience that supervenes upon its phenomenal character, and meanings are not among auditory appearances. This removes one important motivation for the claim that humans hear meanings. My project is not wholly negative. This paper proposes an alternative account of the central phenomenon hearing meanings is supposed to explain: the genuine phenomenal difference in auditory experience that accrues thanks to knowing a language. According to this account, understanding spoken utterances requires learning to hear language-specific sounds.

Sections 1 through 3 explain the motivations and introduce the semantic perceptual account. Section 4 articulates the account's commitments. Section 5 develops the undercutting argument. Sections 6 and 7 present my alternative account of the main explanandum. Section 8 concludes.

1

The claim is that hearing speech uttered in a language you know sometimes involves being auditorily perceptually aware of its meanings or semantic properties. This paper addresses one potentially powerful way to motivate the claim. The motivation begins by recognizing a contrast. The contrast is between the experience of listening to speech in a language you know and the experience of listening to speech in a language you do not know. A vignette illustrates the contrast.

Imagine hearing some speech uttered in a language you know. Now imagine hearing some speech uttered in an unfamiliar language. Since this leaves a lot of leeway, make the example more concrete. Imagine hearing a supermarket cashier talking to you in your native language, maybe telling you that the cookies are on special. Now imagine everything is the same except the supermarket cashier is talking to you in a totally unfamiliar language, telling you that the cookies are on special (you have no idea that is what you are being told). Also imagine explicitly that you are conscious, awake, and aware of what I ask you to imagine in each case. The thing that should strike you as apparent is that the experience of listening to speech in a language you do not know differs from the experience of listening to speech in a language you understand and speak fluently.

Three clarifications help to identify the relevant contrast. First, by *experience* I mean a conscious mental episode with something it is like for you as a subject to undergo. I mean *experience* in the broadest possible sense, so that it may encompass, for instance, sensory, perceptual, bodily, affective, emotional, imaginative, and even occurrent cognitive events or states (experience thus may include what Strawson (2010, 340) calls 'sensory-affective experience' and 'cognitive experience').

In particular, I do not yet wish to limit discussion specifically to auditory or sensory perceptual experiences. The claim is that your overall experience when listening to speech in a familiar language differs from your overall experience when listening to speech in an unfamiliar language.

Second, the specific sort of difference I care about here is a *phenomenal* difference. Call what it is like for you to undergo an experience its *phenomenal character*. Allow that the interesting properties of your experience outstrip what it is like for you as a subject. Your experiences have properties, such as being realized by a physical state or having a certain causal history, that do not enter your consciousness. So your experiences differ in respects that do not enter your consciousness. A phenomenal difference, however, is a difference in phenomenal character, a difference in what it is like for you as a conscious subject to have each experience. There is something it is like for you when you have the experience of listening to the foreign language, and there is something else it is like for you when you have the experience of listening to the familiar language. These differ qualitatively, so your experiences differ phenomenally. That they so differ is in principle noticeable to you just by having and comparing the two experiences.

Third, your experiences on these occasions might have differed phenomenally in many ways. The clerk could have appeared black-haired in one case but redheaded in the other; you could have felt nausea in one case but joy in the other. Forget about those differences. Hold everything fixed but the language spoken. If we control for all the differences that have nothing to do with listening to spoken language, and for accidental details like the average pitch of the clerk's voice or whether the clerk seems to have a frog in the throat, your experiences nonetheless differ in features you could notice if you tried. This is what I mean by saying that the experience of listening to speech in a language you know differs phenomenally from the experience of listening to speech in an unfamiliar language. The argument for this depends on a contrast you can readily detect from the first-person perspective upon having the two experiences in close sequence.

One remaining source of phenomenal difference is that a fluent speaker makes different *sounds* when speaking different languages. Languages use different sound inventories and combine sounds in novel ways. Whether or not you know Japanese, it sounds different from French—the auditory appearance of Japanese speech differs from that of French speech. Hearing different sounds and sound combinations makes the auditory experience and thus the overall experience you have when listening to any one spoken language differ phenomenally from the experience you have when listening to any other.

However, we can control for the acoustical differences among languages that make a phe-

nominal difference to auditory experience by *fixing* the language. What is noteworthy and philosophically significant is that it remains plausible that knowing the language makes a phenomenal difference to the overall experience of listening to speech. In *Mental Reality*, Galen Strawson discusses the phenomenon at length.

[D]oes the difference between Jacques (a monoglot Frenchman) and Jack (a monoglot Englishman), as they listen to the news in French, really consist in the Frenchman's having a different *experience*? ... It is certainly true that Jacques's experience when listening to the news is very different from Jack's. (Strawson 1994, 5–6)

Philosophers whose commitments otherwise diverge agree. Ned Block mentions 'the difference in what it is like to hear sounds in French before and after you have learned the language' (Block 1995, 234), and Michael Tye discusses 'phenomenal differences in what it's like to hear sounds in French before and after the language has been learnt' (Tye 2000, 61). Here are three additional recent examples.

If a monolingual English speaker hears the word *Hund* it is just a sound. If a German speaker hears it, the experience is phenomenally different. (Prinz 2006, 452)

Almost everyone has experienced hearing others speak in a foreign language that one doesn't understand, and that one can't parse into words and sentences. The phenomenology of hearing the same speech when one does understand is markedly different. (Siegel 2006, 490)

There is, intuitively, a difference between what it is like to hear the sentence 'il fait froid' when one does not understand French and what it is like to hear the same sentence after having learnt French, despite the fact that both involve the same auditory input. (Bayne 2009, 390)

Fixing the language makes it tougher to establish that there is a phenomenal difference. The argument now depends on a contrast that cannot be noticed by a single subject from the first-person perspective in a short time span. The contrast holds either between experiences separated by a long time or between experiences of distinct subjects. To convince yourself of the contrast with the language fixed, you cannot just listen in turn to two utterances and compare your experiences. Instead, you have to start out ignorant of the language and then wait a long time until you

have put in the effort to learn it. Alternatively, you have to compare your experience to someone else's experience and figure out if their phenomenal characters differ. This involves some philosophically contentious methodology, but, since it is plausible enough, suppose that it works out. Maybe your memory for phenomenal character is sharp, or you are great at communicating about experiences, or the phenomenal differences are dramatic.

For further support, which requires neither the years nor comparing experiences across subjects, we can appeal to experiences associated with *sinewave speech*. Sinewave speech is an artificial speech signal that is at first very difficult to hear as speech.¹ With help from some prompts, however, it can be grasped as speech and understood. A big phenomenal shift occurs once you start to hear it as comprehensible speech instead of as a bunch of non-linguistic blips and boops. Moreover, there is a discernible but much less pronounced shift when the sinewave speech is in a foreign language you do not understand. Differences in sounds between the languages do not explain the asymmetry between the shifts associated with hearing sinewave speech in known and unknown languages. The difference between the two phenomenal differences is the difference knowing a language makes.

This supports the claim that, controlling for acoustical differences, the experience of listening to an utterance spoken in a language you know differs phenomenally from the experience of listening to an utterance spoken in a language you do not know.

2

What is the source of the difference? A natural proposal is that when you listen to speech in a language you know, you *understand* it. You *grasp* the meanings of the expressions uttered, thanks in part to your syntactic, grammatical, and semantic competence.

This cannot yet support the claim that you *perceive* meanings. Philosophers commonly distinguish perception from thought and other forms of extra-perceptual cognition. For example, seeing an elephant differs from remembering an elephant, and hearing a duck's sound differs from making a logical inference. To avoid confusion about other uses of the term 'perception', I want to emphasize that the target of this paper's discussion is a *sensory* form of perception, or *sensory perception*. By that I mean a variety of perception associated with one or more of the exteroceptive sensory modalities, such as sight, hearing, touch, taste, or smell. If you are sensorily

¹See Remez et al. (1981).

perceiving, you are seeing, hearing, touching, tasting, or smelling. (If there are more or fewer than the traditionally recognized modalities of sensory perception, adjust accordingly.) In particular, the target is *auditory* (sensory) perception, or audition. Moreover, my concern here is with auditory perceptual *experiences* rather than any entirely nonconscious or subpersonal form of audition. Auditory experiences are auditory episodes that have, among other properties, phenomenal character. To avoid prejudging the question whether all experiences have objects or intentional contents, I stipulatively reserve 'awareness' for any mental episode with an object or intentional content. Experiential awareness is any such episode with phenomenal character.

How to distinguish sensory perception from occurrent thought and other varieties of cognition is vexed, and no uncontroversial criterion exhaustively sorts mental episodes into sensory perceptions and thoughts or extra-perceptual cognitions. However, agreement exists about some central cases. Seeing hues and shapes, hearing pitches and loudness of sounds, and smelling pungent odors involve sensory perception. Believing that the square of two is four, recalling your grandmother's advice, and thinking about the possibility of life on Mars involve forms of thought or cognition distinct from sensory perception. Five threads run through diverse philosophical discussions of exteroceptive sensory perception. First, sensory perception frequently involves *vivid* experiences. Second, sensory perception commonly involves awareness *as of* sensible features, which may include qualities, such as hue, pitch, saltiness, and stench; properties, such as shape, volume, and texture; and sensible individuals. Third, in cases of success, sensory perception *depends* upon the features of which you are perceptually aware. Fourth, the features as of which you are sensorily perceptually aware commonly strike you as being *independent* from you and from your sensory perceptions in a way that, for instance, the pain of a headache is not; this helps to distinguish sensory perception from mere *sensation*. Fifth, sensory perception is to a degree *insulated* from influence by what you believe and desire and from your attempts to control it. Each of these threads helps explicate the intuitive idea that, compared to (say) remembering and believing, exteroceptive sensory perception affords a way of getting into a relatively intimate form of conscious mental contact with stuff outside your body.

Linguistic understanding relies not just on hearing, but also on memory, conceptual skills, and a mastery of syntax, grammar, and semantics that is not afforded by sensory perception alone. Grasping meanings thus involves forms of *cognition* distinct from sensory perception (whether grasping meanings involves some non-sensory or wholly cognitive form of perception is beyond this paper's scope). Therefore, the proposal being considered entails that the difference between

the experience of listening to speech in known and in unknown languages includes *cognitive* differences that outstrip auditory perception. Suppose we control for things like mental imagery, associations, emotion, and subvocalization, and focus just on the cognitive grasp of meanings. Accepting the present explanation for any phenomenal difference that remains requires accepting that such cognitive episodes have associated phenomenal properties—that there is something it is like to undergo ‘meaning-experience’ (Strawson 1994, 7). This is the focus of recent debates about whether grasping meanings in thought or cognition has a distinctive proprietary phenomenology.²

Jacques and Jack encounter the same sounds. In the text I omitted by ellipsis in the quotation above, Strawson says, ‘The present claim is simply that Jacques’s experience when listening to the news is utterly different from Jack’s, and that this is so even though there is a sense in which Jacques and Jack have *the same aural experience*’ (6, my emphasis). Suppose Jacques’s auditory perceptual experience is a perfect phenomenal match to Jack’s. If so, the phenomenal difference between the experience of listening to utterances spoken in known and in unknown languages includes no difference in auditory perceptual awareness. It is entirely extra-perceptual.

3

Perhaps awareness of meanings is not entirely cognitive and extra-perceptual. According to some philosophers, listening to spoken utterances in a language you know involves hearing not just sounds but also their *meanings* or *semantic properties*. Concerns about the publicity of meaning lead John McDowell to hold that semantic content is ‘something capable of being heard or seen in the words by those who understand the language’ (McDowell 1998, 99).³ However, one might like this account for a different kind of reason. Plausibly, the phenomenal difference also concerns how you auditorily perceptually experience the speech. The very same speech somehow just *sounds* different when you know the language, in respects that cannot be chalked up entirely to extra-auditory cognition. In a footnote to the statement quoted in the last paragraph, Strawson says, ‘In one sense, of course, they do not have the same aural experience, because of Jacques’s automatic segmenting of the stream of sound into words’ (6). If so, controlling again for acoustical differences, the *auditory perceptual experience* of listening to speech in a known language differs phenomenally from the auditory perceptual experience of listening to speech in an unfamiliar

²See, e.g., Strawson (1994); Siewert (1998); Pitt (2004) in favor; Robinson (2005); Prinz (2006) against.

³See Smith (2009) for explication, commentary, and dissent.

language. This too takes some effort to show. One way to do it is to appeal to qualitative and temporal differences when listening to known and to unknown speech. Imagine you are a monoglot English speaker hearing statements uttered in English, Japanese, and Hindi. Temporally, what sound like exaggerated gaps and pauses between words in the English utterance seem missing from the Japanese and Hindi utterances. Qualitatively, you hear clearly the apparent difference between 'l' and 'r' that a Japanese speaker misses, but miss the difference between aspirated and unaspirated 'p' that a Hindi speaker hears clearly. This topic is complex, and I return to it in §7. The present aim is to motivate thinking that the overall phenomenal difference stems partly from a discernible difference in auditory perceptual experience.

Some philosophers have expressed support for hearing meanings on the grounds that the phenomenal difference between the experience of speech or text in a known language and the experience of speech or text in an unknown language involves a phenomenal difference that is best explained by perceptual awareness of meanings. Call this a *semantic perceptual account* of the phenomenal difference. Tim Bayne says, 'It seems plausible to suppose that the [perceptual] phenomenal differences brought about by learning French involve both (low-level) changes in phonological structure *and* (high-level) semantic differences' (Bayne 2009, 393). Susanna Siegel argues that certain phenomenal contrasts are best explained by the perceptual experience of high-level kind properties, cites as an example that visually experiencing semantic properties of texts best explains the phenomenal contrast between seeing familiar and unfamiliar text, and says that seeing written text is the visual analog of hearing speech (Siegel 2006, 490). When you learn to read, 'You become disposed to attend to the semantic properties of the words in the text, and less disposed to attend visually to the orthographic ones' (490–1). Siegel suggests that viewing highway billboards while quickly passing indicates that you 'take in' semantic properties of their text without deliberately attending to the text, and this 'suggests that the "taking in" [of semantic properties] can be merely sensory' (497). Since seeing text is like hearing speech, when hearing spoken language you auditorily take in meanings or semantic properties in a merely sensory way.

According to a semantic perceptual account, you auditorily experience speech in a language you know to bear meanings or semantic properties, and you do not auditorily experience speech in unknown languages to bear meanings or semantic properties. The motivation for accepting this account is that it explains the salient phenomenal difference between the auditory perceptual experience of listening to spoken utterances in known and in unknown languages.

What are the semantic perceptual account's commitments?

First, as above, we are talking about an auditory form of sensory perceptual experience. One part of the view is that the auditory experience of listening to speech in a language you know involves auditory perceptual awareness as of meanings or semantic properties. Using 'as of', rather than just 'of', allows that in some or all cases you only seem to be perceptually aware of meanings and so that some or all meaning experiences are illusory. As I am using this locution, you can be perceptually aware *as of* something even if you are under an illusion. This fends off the argument that perceptual awareness requires success but there are not any meanings around to be perceptually aware of. Whether or not there are any meanings to hear (treating 'hear' as a success term), the semantic perceptual account entails that you have conscious auditory episodes that are illusory if not appropriately related to meanings. In one popular idiom, the content of your auditory experience includes or is satisfied in part by meanings; in another idiom, you undergo an episode that to you is indistinguishable from being auditorily aware of or acquainted with meanings; in another, you are auditorily aware as of meanings. In what follows, for simplicity, I'll assume that if you ever are auditorily aware as of meanings, you sometimes hear meanings.

Second, it is crucial to recognize that the semantic perceptual account is committed to auditory perceptual awareness as of *specific meanings*, rather than just as of *meaningfulness*. The reason is that it aims to explain the phenomenal difference between auditory perceptual experiences of listening to known and to unknown language, and nothing suggests that you could not hear speech in a foreign language *as meaningful* if hearing speech as meaningful does not require hearing specific meanings. In fact, hearing an utterance as meaningful though not as having a particular meaning plausibly helps explain why hearing and recognizing foreign speech as language differs phenomenally from hearing non-language as such. So you might hear French and Hindi utterances as meaningful language even if you do not hear their meanings. It obviously does not help in this context to object that hearing *meaningfulness* requires hearing specific meanings. However, one might respond by objecting that you can hear an utterance *as meaningful* only when you know the language and thus extra-perceptually grasp its specific meaning. Grasping meanings in that case causally makes a difference to auditory awareness, though you only hear utterances as meaningful if you know the language (and you never hear specific meanings). This could explain the phenomenal difference between auditory perceptual experiences of listening to known and to un-

known speech. The cost is that that is *all* it could explain. Every utterance in every language you know would share a single audible mark of meaningfulness, conferred by understanding it. This, however, is greatly at odds to the view suggested by McDowell's claim that semantic content is 'something capable of being heard or seen in the words,' by Bayne's remark that the relevant phenomenal differences 'involve both (low-level) changes in phonological structure *and* (high-level) semantic differences,' and by Siegel's comment that 'you become disposed to attend to the semantic properties of the words' and 'take in' semantic properties in a 'merely sensory' manner. Taken at face value, the semantic perceptual account entails that auditory perceptual awareness as of specific meanings (rather than just a common mark of meaningfulness) is part of the best explanation for the phenomenal difference between the auditory perceptual experience of listening to utterances spoken in a language you know and of listening to utterances in unfamiliar language.

Third, however, the specific meanings you grasp in understanding an utterance cannot be what you hear. Consider an utterance u . Suppose u is an utterance of an expression that means p . If I understand and grasp the meaning of u , then I grasp p . Grasping p requires thinking, entertaining, or representing p . In particular, understanding u requires entertaining p as an appropriate response to u . But grasping, entertaining, and representing that p do not require endorsing or taking any committal attitude, such as belief or judgment, to p . Consider 'snow is green'. 'Snow is green' means *snow is green*. Understanding 'snow is green' requires being able to think *snow is green*. But I can grasp, entertain, or represent *snow is green* without believing or taking it to be the case that snow is green. While perceiving does not require believing, perceiving is in another respect committal. Imagining aside, I cannot have a perceptual experience in which I perceptually entertain that something is the case, or is present, without having a perceptual experience which purports that it is the case, or is present; so, I cannot perceptually entertain p without being perceptually aware as of p . If I hear the specific meaning of u , then I auditorily grasp, represent, experience, or enjoy awareness as of what I grasp when I grasp the meaning of u . So, if I hear the meaning of u , then I auditorily grasp, represent, experience, or enjoy awareness as of p . However, when I hear u , I do not auditorily grasp, represent, experience, or enjoy awareness as of p . For instance, when I hear an utterance of 'snow is green', I need not auditorily grasp, represent, experience, or enjoy awareness as of snow's being green—I do not hear that snow is green. It does not seem to me auditorily as if snow is green, or as if anything snowy or green is around; my auditory experience is not illusory if snow is not green or if nothing I hear is snowy or green. Therefore, when I hear utterance u , I do not auditorily grasp, represent, experience, or enjoy awareness as of the meaning

of *u*.

This argument does not depend upon utterances of complete sentences, a particular theory of meaning, or construing auditory awareness as a propositional attitude. Given suitable assumptions about perception's relation to thought, versions apply for the full matrix of utterance types, theories of linguistic meaning, and theories of perceptual awareness. If understanding is grasping meanings or semantic contents, then you do not enjoy perceptual awareness as of meanings or semantic contents. Perceptual awareness is committal in a way that understanding is not.

This argument does not refute the semantic perceptual account, but it helps explicate its commitments. The claim cannot be that you auditorily experience *what* you grasp when you grasp *the specific meaning* of an utterance. The claim instead must be that you hear utterances *to have* certain meanings, or hear *that* utterances have certain meanings. When you know the language, hearing an utterance *u* involves your auditorily experiencing that *u* means *p*; or, alternatively, auditorily attributing to *u* the feature *means p*; or, alternatively, being auditorily aware as of *u's meaning p*. The critical point is that you hear features like *means p* rather than meanings (like *p*) themselves.

Suppose this argument is sound. According to the semantic perceptual account, you auditorily experience an utterance to have specific semantic features such as *means p*, or your auditory experience involves awareness that an utterance *means p*. For instance, you hear my utterance of 'snow is green' to mean that snow is green (rather than hearing that snow is green). From now on, this is what I intend when I say that, according to the semantic perceptual account, you hear meanings or semantic properties.

This commitment has a noteworthy consequence. Understanding an utterance does not *require* grasping or entertaining thoughts about the meanings you auditorily experience utterances to bear according to the semantic perceptual account. Thus, the account entails that when you listen with understanding to an utterance, the contents you must grasp in occurrent thought (in order to understand) mismatch your auditorily experience. Consider a young child hearing and understanding an utterance of 'snow is green'. The child grasps what is meant and expressed: *that snow is green*. But suppose the child lacks the concept of *linguistic meaning*. Plausibly, the child is unable to entertain the proposition '*snow is green*' means *that snow is green*. The semantic perceptual account implies, however, that the child has an auditory experience that attributes the property *means that snow is green* to the audible utterance. Some might object that this overintellectualizes auditory awareness. I simply want to call attention to the unexplained asymmetry between auditory awareness and extra-perceptual cognition. You grasp *snow is green* but hear *u means that snow*

is green. Nevertheless, since it frequently is the case that linguistically competent adult listeners acquire beliefs to the effect that *u means p*, or that *s said p*, this again is no refutation. But it makes pressing the need for a strong positive case.

Fourth, the semantic perceptual account entails that there is *something it is like* to hear an utterance to mean *p*. If you auditorily experience semantic features of utterances, then there is something it is like for you to be auditorily aware as of semantic features. This parallels saying that when you visually experience colors (or causation or cows), there is something it is like for you to be visually aware as of colors (or causation or cows).

This suggests the question, What is it like to auditorily experience an utterance to have a specific meaning? A natural response is to describe hearing the acoustical features of the linguistic sounds that correspond to particular meanings. The distinctive audible characteristics of an utterance of the English word 'dog' are the sounds of the spoken word 'dog'. If so, what it's like for a competent English user to auditorily experience an utterance of 'dog' to mean *dog* is just like auditorily experiencing the acoustical features of an utterance of 'dog'. But that cannot distinguish the phenomenal character of an auditory perceptual experience that attributes *means dog* to an utterance of 'dog' from the phenomenal character of an auditory perceptual experience that attributes no meaning to an utterance of 'dog'.

This illustrates that, to explain the phenomenal difference in the auditory perceptual experience of listening to speech in known and in unknown languages, hearing specific meanings must contribute in a *distinctive* manner to the phenomenal character of auditory experience. Auditory awareness as of meanings must affect auditory phenomenology. The claim, however, is not merely that grasping meanings makes a causal difference to the phenomenal character of auditory experience. Suppose you grasp a meaning through extra-perceptual cognition and that this changes how you auditorily experience just the non-semantic features (e.g., pitch, timbre, timing) of the sounds. In that case, there is no phenomenal difference in your auditory experience that cannot be captured in other, non-semantic terms, so grasping meaning causes an auditory phenomenal change without auditory perceptual awareness as of semantic properties. Or suppose that upon learning a language your auditory experiences represent meanings (perhaps for reasons analogous to why you have XYZ thoughts after a long visit to Twin Earth) but that this affects only the phenomenal character of auditory experiences of non-semantic qualities of language sounds. If no phenomenal difference to your auditory perceptual experience cannot be captured in terms of awareness as of non-semantic features, the claim that you auditorily experience meanings is not

required to characterize the phenomenal difference. Meaning experience needs to contribute, and in a constitutive rather than merely causal way, to the phenomenal character of auditory perceptual awareness.

What, then, is phenomenally distinctive about the auditory perceptual experience of specific meanings or semantic properties? One way to answer would be to say how such features auditorily seem to be, or to say how sounds or spoken utterances seem to be just when they are heard as having meanings. This is contentious because it relies on detailed descriptive reports based on introspection. Moreover, auditory awareness as of meanings might affect auditory experience in a distinctive phenomenal manner even though you do not auditorily experience specific meanings *to be* any introspectively identifiable way. Perhaps you hear an utterance as having a given meaning without recognizing the semantic property as such; perhaps you experience meanings as audible similarities and differences among utterances; or perhaps the complexity of experience makes it immensely difficult to isolate meaning experiences. A full characterization of what it's like for you as a conscious subject to hear utterances in a familiar language nevertheless might require mentioning auditory awareness as of meanings. However, even if the semantic perceptual account need not offer a detailed phenomenological description of the auditory experience of meanings, it nonetheless does require an argument that there are distinctive phenomenal markers associated with hearing the specific meanings of utterances.

5

The previous section posed the question, What is it like to hear an utterance as having a specific meaning or semantic property? It perhaps sets the bar too high to ask for an informative descriptive answer. But even if we cannot say in detail what it is like to auditorily experience specific meanings, the semantic perceptual account holds that hearing specific meanings makes a phenomenal difference. The account requires that auditorily attributing specific meanings to utterances shapes auditory experience in a distinctive phenomenal respect.

Does auditory awareness as of specific meanings make a phenomenal difference? The way to demonstrate that hearing specific meanings makes a distinctive phenomenal difference is to identify a clear phenomenal difference in auditory experience that traces to a difference in meaning and that lacks a better explanation. A straightforward and direct case should involve two auditory experiences in which the audible non-semantic features match but in which a difference in meanings

makes a discernible phenomenal difference. We want the same sounds, different meanings, and a phenomenal difference in auditory perceptual experience.

This may sound puzzling. You may be thinking that the featured example of this paper is just such a case. But it is not. The featured example concerns the phenomenal contrast between listening to speech in a foreign language and listening to speech in a familiar language. That contrast shows that sometimes when you grasp the meanings of spoken utterances your auditory perceptual experience differs phenomenally from when you do not grasp any meanings. But it cannot show that *which* particular meanings you grasp makes any phenomenal difference. So it could only show that grasping some meaning or other, rather than no meaning at all, makes a phenomenal difference to auditory experience. A compelling argument that humans auditorily experience specific meanings requires showing that audible differences in meaning affect the phenomenal character of auditory perceptual experience. That requires a different kind of contrast argument.

What's needed is a case that involves hearing utterances that share sounds but differ in meaning. *Homophones* are expressions that differ in meaning but not pronunciation, and they provide just such cases. Homophones include homonyms, which share a spelling, as with 'bank' and 'bank', and heterographs, which do not, as with 'to', 'too', and 'two'.

Consider hearing someone utter homophones. For instance, imagine me meaningfully uttering 'pole', 'pole', and 'poll', in sequence. My three utterances involve acoustically identical sounds. Do you detect any audible differences? Do your auditory perceptual experiences differ in phenomenal character? In this case, it is far more plausible to say that there are no audible differences and that your auditory perceptual experiences match in phenomenal character than it is to say that there is either sort of difference. My utterances' meanings differ, but your auditory experiences do not differ phenomenally.

Maybe in the case just described you hear *no* meanings. Since you are not sure which word is which, perhaps when you hear homophones uttered in sequence you hear no meanings. If you ever auditorily ascribe semantic properties, however, it is plausible that you hear my meaning when I utter 'otter'. So you should hear my meaning when I utter 'poll' once. Since nothing obvious breaks down when I add additional homophonic utterances, you should hear multiple meaningful utterances. Alternatively, then, you might hear *the same* meaning three times in a row.

So revise the example. Suppose now that I tell you the first word I utter means *a long slender piece of rigid material*, the second means *a place on the surface of Earth at its axis of rotation*, and the third means *a vote*. And suppose when I utter them again you clearly grasp or associate each spe-

cific meaning with its word. Still, it is most plausible that no detectable auditory difference exists. If, with your full knowledge, we swapped the words around, you'd likely report no changes in what it's like to auditorily experience the sequence. This is a case in which you understand the meanings of words; those meanings differ; you grasp the different meanings upon hearing the utterances; and no phenomenal difference holds among your auditory experiences. We control for acoustical features and alter the meanings of utterances, but it makes no clearly discernible phenomenal difference to your auditory perceptual experience.

One good objection is that single word utterances do not reveal the phenomenal difference that auditory perceptual awareness as of specific meanings makes because hearing meanings requires expressions embedded in appropriate linguistic contexts. Consider three sentences that informatively embed 'pole', 'pole', and 'poll', respectively. For instance, 'Ernest used the pole to vault over the high bar', 'Last year Mac visited the southern pole of Earth', and 'Bubb won the greatest number of votes in our latest poll'. The claim must be that furnishing semantic context by providing the whole sentence makes a phenomenal difference to your auditory perceptual experience of hearing an utterance of 'pole', 'pole', or 'poll' that is due to your hearing meaning. Ask a friend to utter the three sentences aloud while you try to detect whether hearing the relevant bit differs. Are you able to attend to anything audible in the three utterances of the 'pōl' sound in virtue of which they noticeably differ? Can you detect anything that plausibly counts as an auditory phenomenal contrast while hearing the 'pōl' sounds? While proponents of the semantic perceptual account should accept that supplying linguistic context does generate such auditory phenomenal differences, my own strong inclination is to answer 'no' to each question. One reason beyond introspection is that it implies a mismatch between the conditions under which you grasp semantic properties (which doesn't require such embedding) and those under which you auditorily attribute the corresponding semantic properties. Another reason is that even failing to understand the relevant word seems to leave auditory experience phenomenally intact. For instance, suppose for the moment that you do not know the meaning of 'sedulous', and I utter, 'If McShane weren't so sedulous, Ezra would be worse off'. If you're listening, even if you didn't understand the word 'sedulous', you can repeat its sound and recognize it in another sentence. Once I inform you that it means *diligent* and *careful*, does your auditory perceptual experience change? Suppose instead that I tell you it means something entirely different, such as *wealthy*. If you auditorily experience specific meanings, this should impact the uttered word's audible appearance. But neither learning nor redefining a word's meaning seems to impact auditory appearances or to

be reflected in auditory capacities, so nothing suggests supplying disambiguating context makes the crucial difference.

Adding linguistic context can introduce audible differences beyond the meanings of homophonic expressions. This is a source of potentially confounding phenomenal judgments, so control for it. Note that adding linguistic context by itself may not suffice to resolve audible semantic differences. The auditory experience of hearing, 'Where is the pōl?' does not clearly differ phenomenally under its differing interpretations.⁴ Homophones are not limited to single words. So we need to allow that broader context sometimes is required to hear specific meanings and thus impacts the phenomenal character of auditory perceptual experience. Granting this, listening to homophonic complex expressions and complete sentences should lead to auditory experiences that differ phenomenally. Consider structurally ambiguous statements, such as, 'Visiting relatives can be boring', or those with scope ambiguities, such as, 'Everyone loves someone'. Or consider *oronyms*, such as 'mint spy' and 'mince pie', or 'Some others I've seen were plaid' and 'Some mothers I've seen wear plaid'. Does the experience of listening to such an utterance differ phenomenally in auditory respects depending on the interpretation? This is trickier. In some cases, such as oronyms, there is room to think it does, but that plausibly has to do with how the sounds are segmented and grouped. In other cases, such as scope ambiguity, no clear auditory phenomenal difference exists. Even when the broader context makes clear the meanings expressed, such complex utterance pairs may strike one as audibly indistinguishable at the pace of normal speech.

A final objection is that considerations from homophony aren't fair because we auditorily experience homophones only to have indeterminate or disjunctive meanings or semantic properties. Hearing indeterminate or disjunctive meanings would cancel any phenomenal differences. My response is that this defangs the claim that humans hear meanings. According to the objection, the only meaning differences that make a phenomenal difference to auditory perceptual experience are ones to which there corresponds a perceptible acoustical or non-semantic difference. On such an account, hearing specific meanings makes no phenomenal difference to auditory perceptual experience that is not marked by an apparent acoustical or non-semantic difference. Thus, there could be no compelling argument based on phenomenal contrast that an audible difference in meaning makes a phenomenal difference to auditory perceptual experience.

The arguments from homophony cast doubt on the claim that semantic differences make any phenomenal difference to auditory perceptual experience in absence of other audible differences.

⁴Siewert (1998, ch 8) discusses an example of this kind in relation to cognition.

One might simply insist that there is an auditory phenomenal difference in the cases described above, but that claim is difficult to motivate without a compelling contrast scenario. Homophones show that the phenomenal character of auditory experience does not settle questions of meaning in the straightforward way that it settles questions of pitch, loudness, duration, and even whether an utterance is a speech-like sound. Thus, the homophone argument for hearing specific meanings is unconvincing.

6

The central argument being considered for the claim that you hear meanings is that it explains the phenomenal difference between the auditory perceptual experience of listening to speech in a known language and the auditory perceptual experience of listening to speech in an unfamiliar language. According to the semantic perceptual account, the auditory experience of specific meanings of utterances in a language you know affects phenomenal character in a distinctive manner. The previous section, however, argued that since homophones ground no convincing phenomenal contrast argument, grasping different meanings makes no clear phenomenal difference to your auditory perceptual experience. This argument undermines support for the semantic perceptual account by showing that a consequence of the commitments outlined in §4 fails. But that threatens to leave us with no explanation for the phenomenal difference. One option is to deny that knowing a language affects the phenomenal character of auditory perceptual experience and to argue that any phenomenal difference stems from extra-perceptual cognition. This, however, backtracks and rejects the plausible assumption that motivated the claim that you hear meanings. Since the central argument for hearing meanings is that it explains a phenomenal difference in auditory perceptual experience, I prefer to grant the premise, provide an alternative explanation for the phenomenal difference, and argue that it is a better explanation than hearing meanings.

We can give a good account of the phenomenal difference knowing a language makes to the auditory perceptual experience of speech without appealing to the perception of meanings or semantic properties. How? The *prima facie* obstacle is that ignorance of language implies no auditory deficits. In my examples, your ears are fine and conditions are good for hearing the sounds of the unfamiliar language. What could change auditorily once you learn it? The way forward is to recognize that learning a language involves mastering more than just its sound-meaning mapping. What the semantic perceptual account gets right is that, prior to learning the language, you

in fact are at a kind of auditory disadvantage; your auditory experience is impoverished relative to those who know the language. But this isn't about failing to hear meanings. Language learning also requires learning to *hear* the sounds of the language, in a sense I'll explain. My strategy is to account for the phenomenal difference in auditory experience by appealing to differences associated with hearing the sounds of the language rather than its meanings. The evidence favors this account of the perceptual phenomenal difference over the semantic perceptual account. Grasping meanings nonetheless may contribute to the phenomenal character of extra-perceptual cognition.

7

Learning a language alters the temporal and qualitative features speech sounds audibly appear to have. For example, hearing foreign language is like hearing a mostly unbroken sound stream. Speech in your native language, however, auditorily seems segmented into units separated by gaps and pauses. Barry C. Smith describes listening to speech in a foreign language in these terms: 'In such circumstances, one hears not words but a continuous sound stream interrupted when the speaker pauses for an intake of breath' (Smith 2009, 185). I quoted Strawson earlier as saying, 'In one sense, of course, they do not have the same aural experience, because of Jacques's automatic segmenting of the stream of sound into words.' Hearing known language also involves distinguishing sublexical features at a finer temporal grain and making out subtler qualitative differences. In principle, we might be able to explain the attendant phenomenal differences in terms of changes in patterns of auditory awareness as of relatively low-level audible qualities, such as pitch, timbre, and loudness.

Such a conservative account of the experiential difference need not suffice. Learning a spoken language involves gaining auditory perceptual awareness as of new kinds of features—in particular, non-semantic but *language-specific* attributes of spoken utterances. This awareness affects the phenomenal character of auditory perceptual experience. So we can offer an even richer account.

Phonological features are one plausible though not uncontroversial candidate, so I'll focus on them. To illustrate, first consider utterances of 'top', 'ute', and 'rat'. These utterances in one respect are audibly similar. Each audibly shares with the others the 't' sound—[t], in phonological notation. Next, consider the audible difference between utterances of 'tan' and 'ran'. These audibly match, except that one includes the 't' sound where the other includes the 'r' sound—[ɹ], in phonological notation ([r] is the trill sound). The [t] and [ɹ] sounds, or *phones*, are examples of

audible features that are shared among distinct utterances and that differ among otherwise indistinguishable utterances. Since they are the smallest such features that make a linguistic difference in *any* of the world's languages, they are the audible 'building blocks', interpreted very loosely, of spoken language. By using 'building blocks', I am alluding to but not endorsing the 'beads on a string' analogy.⁵ I do not endorse that analogy because I hold that the sound features I am talking about are *properties* or attributes (of utterances) rather than individuals, and that their instances are *overlapping* rather than discrete. It is useful to think of phones in terms of minimal equivalence classes of sounds that make a common linguistic contribution to or share linguistic significance in some world language.

Distinct spoken languages, however, do not make use of the same building blocks, nor do individual languages yield the same basic equivalence classes of linguistically significant sounds. I do not just mean that only certain spoken languages include clicks and buzzes. Even when the basic types of sounds they utilize overlap, spoken languages nonetheless differ in which sounds they treat as linguistically equivalent and in which sounds they treat as linguistically distinct. For instance, spoken English distinguishes [l] and [ɭ], but Japanese does not. Thus, while the *phones* [l] and [ɭ] correspond to the distinct English *phonemes* /l/ and /ɭ/, [l] and [ɭ] are *allophones* or linguistically-equivalent variants of a single Japanese phoneme. Thus, the difference between [l] and [ɭ] suffices to distinguish English but not Japanese words. Mandarin Chinese distinguishes [p] and [p^h] as distinct phonemes, but English treats them as allophones of the single phoneme /p/. Many more examples from phonology support the same basic point. Some languages treat certain sound pairs as linguistically equivalent while other languages treat them as distinct. So a spoken language's basic stock of sounds is distinctive to that language.

The really telling explanatory point is that in many cases humans who know one language hear certain crucial pairs of utterances as audibly equivalent, while humans who know another language hear them as distinct. To take one sort of example, auditory perceptual discrimination tasks in linguistic contexts indicate that the sounds that correspond to 't' in utterances of 'ton' and 'stun' auditorily appear virtually the same to users of English or French but differ noticeably to users of Chinese. Utterances of 'bed' and 'bad' audibly differ to English users but not to Dutch users. In some such cases, the acoustical signals are interchangeable without detection. Examples of this sort multiply as above. (Another sort of example is that users of one language

⁵The controversial 'beads on a string' analogy traces to Bloomfield (1933) and is related to the 'alphabetic conception' criticized by Appelbaum (1999).

sometimes hear a shared linguistic sound across acoustically varied utterances while non-users do not.) So suppose we have one group of attentive language listeners who, when presented with two linguistic sounds, systematically notice no difference between the sounds, behave as if the sounds are audibly equivalent, and judge that the sounds are audibly equivalent. And suppose that we have another group of attentive language listeners who, when presented with the same two linguistic sounds, systematically notice an audible difference, behave as if the sounds differ audibly, and judge that the sounds are not audibly equivalent. (A parallel argument using similarities stems from the previous parenthetical example.) It is reasonable to hold that, for at least one of the linguistic sounds, the auditory perceptual experience of an attentive language listener from the first group differs phenomenally from the auditory perceptual experience of an attentive language listener from the second group. So, for a broad class of linguistic sounds, there are noteworthy phenomenal differences between the auditory perceptual experiences of subjects who know different languages. Hearing language-specific features, such as phonemes, elegantly explains such patterns of difference. For example, having auditory perceptual experiences that involve awareness as of the English phonemes /l/ and /ɹ/ accounts for the phenomenal difference between hearing utterances of 'raw' and 'law'. Having auditory perceptual experiences that involve awareness as of a single Japanese phoneme accounts for the lack of phenomenal difference between hearing utterances of 'raw' and 'law'. Thus, an important part of the phenomenal difference associated with listening to spoken utterances in a language you know stems from auditorily experiencing attributes whose linguistic significance is specific that language.

Phonemic differences, unlike semantic differences, affect the phenomenal character of auditory experience even without acoustical differences. Embedding a given acoustical signal in different contexts commonly leads to phoneme experiences that differ phenomenally. Here are four examples.⁶ The acoustical signal that corresponds to /p/ in an utterance of 'pi' is nearly identical to that which corresponds to /k/ in an utterance of 'ka'. A complex of low-level audible qualities that sounds unequivocally like /d/ in one context can sound clearly like /t/ in another, especially from the mouths of different speakers or surrounded by different phonemes. A puff of air on your neck, when presented with a stimulus midway between /b/ and /p/, makes you hear /p/ and alters the phenomenal character of auditory experience. In the McGurk effect, seeing lips move as in /ga/ dramatically changes the phenomenal character of auditory perceptual experience and shifts an otherwise audibly apparent /ba/ to /da/. These considerations block a

⁶See Cooper et al. (1952); Gick and Derrick (2009); McGurk and MacDonald (1976).

counter-argument parallel to that of §5.

Additional evidence supports the following claims: understanding heard spoken language requires auditorily perceiving language-specific attributes, such as phonological features, that are not meanings; auditorily perceiving such non-semantic language-specific features affects the phenomenal character of auditory perceptual experience; and auditorily perceiving such language-specific features requires substantial exposure to the language. Thus, humans learn to hear non-semantic language-specific features of spoken utterances.

First, it is widely accepted that *developmental* evidence shows that young infants are capable of distinguishing all of the phonetic differences that are linguistically significant to any world language. However, between five and twelve months, infants lose the ability auditorily to discern differences that are not linguistically significant for their language(s). American babies stop distinguishing utterances of [k] and [k^h], and Spanish babies stop distinguishing [s] from [z]. And they do this prior to learning the meanings of words. This sort of pruning, which requires experience and sustained exposure, is part of learning to hear the distinctive audible features of a spoken language.⁷

Second, adults auditorily perceive many language-specific speech sounds, such as stop consonants, *categorically*. No clear, univocal conception of what it means to perceive categorically is shared among researchers, but here is the most relevant evidence.⁸ Audibly apparent relations of similarity and difference among language-specific speech sounds do not map neatly onto acoustical similarity and difference; in particular, gradual variation along certain acoustical dimensions produces sharp or abrupt auditory differences. So two token speech sounds with very similar auditory appearances and that seem auditorily to be utterances of the same phoneme might be just as different on a diagnostic acoustical dimension (e.g., voice onset time) as two token sounds with very different auditory appearances and that seem auditorily to be utterances of different phonemes. Detecting an acoustical difference of a given magnitude thus is much easier between utterances that belong to distinct language-specific categories (such as the phonemes /b/ and /p/) than between utterances that belong to the same language-specific category. For acoustical measures that are diagnostic of whether something is a /b/ or /p/, for instance, acoustical differences within a category make little audibly apparent difference, while acoustical differences

⁷Eimas et al. (1971) is the seminal paper; Jusczyk (1997) provides a good start in this literature.

⁸Harnad (1987) is a classic resource on categorical perception, and contains a number of informative discussions of categorical perception and learning for speech. Harnad (2003) is a nice, brief summary.

that span categories make big audibly apparent differences. The upshot is that categorical perceptual processes yield distinctive patterns of audibly apparent similarity and difference among speech sounds. Categorical perception is shaped by learning, however, and these patterns vary by language. Perceiving language-specific features of utterances therefore not only explains characteristic patterns of audibly apparent similarity and difference among speech sounds for users of a given language but also explains why these patterns vary across languages.

Third, further evidence from *aphasias*, language-related disorders, suggests that understanding spoken language requires hearing language-specific features of speech that are not meanings. Patients with transcortical sensory aphasia (TSA) have serious difficulty grasping meanings. They produce fluent speech, but it is jumbled and semantically very confused (paraphrastic). Still, they can repeat back what is spoken to them, and they often have little idea they suffer from the disability. Contrast this with pure word deafness (PWD). Those with PWD produce fluent, meaningful speech and understand written and signed language. However, individuals with PWD cannot auditorily grasp spoken language (unless they've learned to lip-read). Their general hearing is good, and they can hear cars, clapping, and barks. They even can recognize familiar voices. While some identify spoken language as such, none can make out the words or repeat the language-specific sound sequences uttered. Individuals with PWD describe the auditory experience of spoken language as like hearing noise, garbled sound, or foreign language—speech sounds do not 'register' or 'come up'.⁹ Suppose on the basis of these descriptions that PWD shows that the auditory experience of speech when one hears and recognizes the words differs phenomenally from the auditory experience of speech when one cannot. TSA, however, shows that listeners can hear the words and sentences uttered without grasping their meanings (and without noticing the trouble). So there is a salient phenomenal difference between the auditory experience of speech when one recognizes the words and when one does not recognize the words that does not depend on grasping meanings. Auditory perceptual awareness as of non-semantic, language-specific features, such as phonemic and other structural properties of utterances, plausibly captures the phenomenal difference.

We now have a promising case that listening to a language you know involves auditorily perceiving language-specific attributes of utterances rather than just the audible qualities you hear non-linguistic sounds to have. Moreover, perceiving such features affects the phenomenal character of auditory perceptual experience. The case of phonological features is telling, but it is not particularly special. There are other candidates for higher-level structural features you perceive

⁹See the descriptions and illuminating quotations in Poeppel (2001, esp. 681).

spoken language to bear. Some include syllables, phonologically-individuated words, prosody, syntactic features, and perhaps even grammar. Each case should be approached on its own terms. Richer linguistic structure and more abstract linguistic properties tend to favor thinking that the relevant awareness involves extra-perceptual cognition rather than sensory perception. Perhaps perception and cognition occupy a spectrum without well-defined boundaries; if so, the case of listening to spoken language is a revealing study.

The preceding discussion argues for three claims. First, understanding spoken language by listening requires auditory awareness as of language-specific features that are not meanings. Second, such awareness affects the phenomenal character of auditory perceptual experience. Third, hearing such language-specific features requires a history of engagement with the language. Thus, you cannot auditorily experience all of the language-specific features required to understand spoken utterances in an unfamiliar language, and the phenomenal character of your auditory experience reflects this. It follows that the auditory perceptual experience of non-semantic language-specific features is part of the best explanation for the phenomenal difference between the auditory perceptual experience of listening to spoken utterances in known and in unknown languages.

8

I've argued that any adequate explanation for the phenomenal difference between the auditory perceptual experience of listening to speech in known and in unfamiliar languages must appeal to two sorts of difference. Knowing a language impacts the audibly apparent lower-level qualitative and temporal features of spoken utterances. And the language-specific features audition ascribes to utterances in an unfamiliar language are impoverished relative to utterances in a known language. Learning a spoken language in these respects involves learning to hear its sounds. However, the argument from homophony shows that meaning differences without audible non-semantic differences do not affect the phenomenal character of auditory perceptual experience. Thus, the auditory experience of specific meanings needn't figure in our account. Since this is compatible with holding that grasping meanings affects the phenomenal character of occurrent extra-perceptual cognition, it is hard to argue that an account of this shape leaves out anything important. It is a good alternative.

Hearing meanings thus is not part of the best explanation for the phenomenal difference knowing a language makes to the auditory perceptual experience of listening to spoken utterances. But

now we lack an argument that auditory awareness as of specific meanings makes a distinctive contribution to the phenomenal character of auditory perceptual experience. So that removes the motivation for accepting that humans hear meanings that is the subject of this paper. If meanings are among the contents of auditory perceptual experience, they are not among those that constitutively shape or that supervene upon its phenomenal character. In this sense, we get by without hearing meanings.¹⁰

Rice University

References

- Appelbaum, I. (1999). The dogma of isomorphism: A case study from speech perception. *Philosophy of Science*, 66:S250–S259.
- Bayne, T. (2009). Perception and the reach of phenomenal content. *The Philosophical Quarterly*, 59(236):385–404.
- Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18:227–247.
- Bloomfield, L. (1933). *Language*. Holt, New York.
- Cooper, F. S., Delattre, P. C., Liberman, A. M., Borst, J. M., and Gerstman, L. J. (1952). Some experiments on the perception of synthetic speech sounds. *Journal of the Acoustical Society of America*, 24:597–606.
- Eimas, P. D., Siqueland, E. R., Jusczyk, P., and Vigorito, J. (1971). Speech perception in infants. *Science*, 171(3968):303–306.
- Gick, B. and Derrick, D. (2009). Aero-tactile integration in speech perception. *Nature*, 462(7272):502–504.
- Harnad, S. (1987). *Categorical Perception: The Groundwork of Cognition*. Cambridge University Press, Cambridge, UK.

¹⁰Many thanks to Steve Butterfill, Alex Byrne, Philippe Chuard, Justin Fisher, Matt Nudds, Christopher Peacocke, George Sher, Barry C. Smith, and Cara Spencer; to audiences at Oxford and at Rice; and to an anonymous referee, for helpful conversations and comments that shaped the development of this paper. Jeff Speaks deserves special thanks for invaluable conversations, exchanges, suggestions, and written responses.

- Harnad, S. (2003). Categorical perception. In *Encyclopedia of Cognitive Science*. Nature Publishing Group/Macmillan.
- Jusczyk, P. W. (1997). *The Discovery of Spoken Language*. MIT Press, Cambridge, MA.
- McDowell, J. (1998). *Meaning, Knowledge, and Reality*. Harvard University Press, Cambridge, MA.
- McGurk, H. and MacDonald, J. (1976). Hearing lips and seeing voices. *Nature*, 264:746–748.
- Pitt, D. (2004). The phenomenology of cognition, or, *What is it like to think that p?* *Philosophy and Phenomenological Research*, 69(1):1–36.
- Poeppl, D. (2001). Pure word deafness and the bilateral processing of the speech code. *Cognitive Science*, 25:679–693.
- Prinz, J. (2006). Beyond appearances: The content of sensation and perception. In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*, pages 434–459. Oxford University Press, Oxford.
- Remez, R. E., Rubin, P. E., Pisoni, D. B., and Carell, T. D. (1981). Speech perception without traditional speech cues. *Science*, (212):947–950.
- Robinson, W. S. (2005). Thoughts without distinctive non-imagistic phenomenology. *Philosophy and Phenomenological Research*, LXX(3):534–561.
- Siegel, S. (2006). Which properties are represented in perception? In Gendler, T. S. and Hawthorne, J., editors, *Perceptual Experience*, pages 481–503. Oxford University Press, New York.
- Siewert, C. (1998). *The Significance of Consciousness*. Princeton University Press, Princeton.
- Smith, B. (2009). Speech sounds and the direct meeting of minds. In Nudds, M. and O’Callaghan, C., editors, *Sounds and Perception: New Philosophical Essays*, chapter 9, pages 183–210. Oxford University Press, Oxford.
- Strawson, G. (2010/1994). *Mental Reality*. MIT Press, Cambridge, MA, 2nd edition.
- Tye, M. (2000). *Consciousness, Color, and Content*. MIT Press, Cambridge, MA.