

PITCH

CASEY O'CALLAGHAN

1 Introduction

Some sounds have pitch, some do not. A tuba's notes are lower pitched than a flute's, but the fuzz from an untuned radio has no discernible pitch. Pitch is an attribute in virtue of which sounds that possess it can be ordered from "low" to "high". Given how audition works, physics has taught us that frequency determines what pitch a sound auditorily appears to have.

A theory of sounds must address the audible qualities. What exactly is pitch, and how is it related to the physical property of frequency? Given psychophysical evidence which indicates that there is neither a linear nor a logarithmic function from the frequency of a sound to the pitch listeners experience it to have, auditory researchers have adopted the view that pitch is a subjective or psychological property. That is, pitches strictly only belong to experiences. I shall suggest, however, that pitch—though not identical with frequency—can be identified with a physical property of sounds themselves. The Standard (Subjectivist) View of pitch does not follow from the failures of *simple* forms of physicalism about pitch; indeed, a promising physical candidate for pitch can be extracted from recent auditory research on pitch perception. According to this Alternative (Physicalist) View, pitch lacks the independent scientific interest warranted by the Simple Views; it is, however, of interest from an anthropocentric point of view. Pitch is nonetheless an objective physical property of sounds.

2 Simple Views of Pitch

2.1 The Very Simple View

The *Very Simple View* of pitch can be extracted from the story we learned about sounds in primary school. We were taught that the pitch of a sound is the frequency of a pressure wave that travels through a medium such as air, water, or helium.¹ Whatever doubts we have about the natures of colors, accepting the Very Simple View makes us happy to say that pitches reside in the world beyond our minds. That is, that the quality we discern when we hear pitch is just the property of having some frequency. When we tune a trombone, one way to get its pitch right is by adjusting the frequency of an F with an electronic tuner.

This Very Simple View accounts for salient aspects of the experience of pitch. It explains the linear ordering of pitches: as frequency increases, so does perceived pitch. A natural account of the musical relations also follows from the Very Simple View of pitch. Small whole-number frequency ratios form the bases of the octave (1:2), fifth (2:3), fourth (3:4), et al. One gets the palpable sense that the natures of such audible relations are revealed by this discovery.

2.2 The Simple View

The Very Simple View is far too simple. The identification of pitch with frequency approaches adequacy for *pure* or *sinusoidal* tones—sounds whose accompanying waves are constituted by sinusoidal pressure variations.² (See Figure 1: Sinusoidal Motion, p. 26). For sinusoidal tones, increasing frequency increases perceived pitch, and decreasing frequency decreases perceived pitch.

Pure sinusoidal tones are rarely encountered in nature, and many complex sounds that are not themselves sinusoids are perceived to have pitch. Fourier showed that any complex sound which is not itself a sinusoid can be analyzed

¹Frequency is just the number of cycles per second, measured in *Hertz* (Hz).

²I shall argue in §3, however, that the Very Simple View fails even for pure sinusoidal tones.

in terms of sinusoids of various frequencies in differing proportions. Fourier's Theorem applies in virtue of the additive principles of wave-like motion. So we can characterize any sound by citing the (maximum or root-mean-square³) amplitude or intensity of each of its sinusoidal constituents. (See Figure 2: Fourier Composition of a Square Wave, p. 27).

The constitutive role of individual sinusoids can be taken quite literally: the phenomenon of *sympathetic resonance* demonstrates that complex tones are *made up* of sinusoids of different frequencies. A pure sinusoidal tone from one tuning fork causes a second tuning fork with the same characteristic frequency to begin sounding in virtue of its own induced sinusoidal motion. A complex tone from a human voice, one of whose Fourier-predicted components shares a tuning fork's characteristic frequency, induces that tuning fork to resonate just as a sinusoidal tone of that frequency does. But a complex sound without a Fourier component at the tuning fork's characteristic frequency induces at best a weakened resonant sounding. The best explanation of sympathetic resonance is that a tuning fork's resonant sounding is caused by a sinusoidal constituent which is genuinely present in the complex sound.

Pure sinusoidal tones are a variety of *periodic* sound. That is, they repeat a certain motion regularly over any given interval. Some complex sounds are also periodic. As a matter of empirical fact, just the periodic sounds have pitch. Though periodic tones may occur within more "messy" or "noisy" sources and thus cause pitch experiences, and sounds that are not exactly periodic may appear to have pitch, pitched tones are generally periodic sounds. If a complex signal is periodic and repeats at regular intervals, the frequency of each of its components (or *partials*) must have an integer-multiple relationship to a *fundamental* frequency. The fundamental frequency of a complex sound is the greatest common whole-number factor of the sound's constitutive frequencies. Thus, the fundamental frequency of the complex (square) tone in Figure 2 is 1000 Hz.

³The *root-mean-square* amplitude is a measure of the average magnitude of a varying instantaneous amplitude. It is given by the square root of the mean of the squared instantaneous amplitudes. For a sinusoid, the root-mean-square value is 0.707 times the peak amplitude. Gelfand (1998), pp. 19–21.

Helmholtz demonstrated that the fundamental frequency of a complex periodic tone determines its perceived pitch.⁴ So the *Simple View* of pitch is that the pitch of a sound is identical with its fundamental frequency; that is, the pitch of a periodic sound is the greatest whole-number frequency by which the frequency of each of its sinusoidal components is divisible without remainder. According to this formulation a sinusoidal component at the fundamental frequency may or *may not* actually be present in the complex sound. The phenomenon of the “missing fundamental” demonstrates that the fundamental frequency component need not be present for a sound to have the same pitch as a sinusoid of that frequency. A tone with constituents at 1200 Hz, 1300 Hz, and 1400 Hz has the same perceived pitch as a 100 Hz sinusoid.⁵ Telephones, which filter low frequencies, illustrate the principle. One hears a man’s voice to have the same pitch in person and over the telephone, though the fundamental is absent from the telephone speaker’s sound.⁶

The Simple View identifies pitch with fundamental frequency, which depends upon the frequencies of a sound’s sinusoidal constituents. It follows that pure sinusoidal sounds and complex sounds such as those that exhibit sawtooth or square wave patterns can share the same pitch in virtue of shar-

⁴Helmholtz (1954).

⁵See Helmholtz (1954) and Schouten (1940).

⁶Terhardt, e.g., (1974) and (1979), has argued that this simple conception is inadequate to account for all types of pitch phenomena. He has distinguished between *spectral* pitch, which is pitch heard in virtue of constituent frequencies that are actually present in the sound, and *virtual* pitch, which is heard despite the absence of a spectral constituent. Terhardt has pointed out that a view analogous to the Simple View fails to account for the fact that a single complex sound is often heard to have multiple pitches, determined by both spectral and virtual pitches. He has also noted that *both* types of pitch may be heard simultaneously, even at the same frequency! Terhardt’s conception depends, however, upon individual pitches, spectral and virtual, being determined by constitutive frequencies that are present in the sound. Indeed, individual spectral and virtual pitches are identified with particular frequencies. So the Simple View that pitch is fundamental frequency may not accommodate all salient pitch phenomena, but it is sufficient for my purposes that once determined, pitches are identifiable with particular frequencies, and that pitch determination is a matter that depends entirely upon which frequencies are present in a complex sound.

ing a fundamental frequency. The Simple View identifies pitches with particular frequencies, and thus preserves the Very Simple View's influential account of the musical relations.

It has the further advantage of being supported by the physiology of auditory perception. The basilar membrane, located within the cochlea, is like a long trapezoidal ribbon, different parts of which most actively resonate in sympathy with a particular sinusoidal frequency. The basilar membrane thus performs a sort of Fourier analysis, decomposing a complex signal into its sinusoidal components. This spectral information is converted into electrical potentials by the hair cells, which activate auditory neurons. Given frequencies activate portions of the auditory nerve that are “tuned” to those frequencies, and this tuned or *tonotopic* organization continues up through the auditory cortex where higher cognitive processes determine pitch or fundamental frequency from spectral information about the sound.⁷ The significance of Fourier decomposition and tonotopic organization is that frequency parameters appear to be preserved and represented through various stages of the auditory physiology. In fact, electrodes can recover and reproduce a sound presented to the ear from the subsequent auditory nerve signal.⁸ It appears that pitch perception depends upon determining the frequencies present in a complex signal. If pitch is fundamental frequency, the problem of pitch perception is physiologically tractable.

3 Problems with the Simple View

The results of psychoacoustics give strong evidence of a subjective pitch scale that correlates rather loosely with fundamental frequency. Though pitch changes only if frequency changes, the magnitude of a pitch change is neither identical with nor a constant function of the magnitude of its corresponding

⁷A cat's cortex is arranged in individual columns which are tuned to characteristic frequencies. See Woolsey (1960) and Gelfand (1998), Chapter 6.

⁸This is the so-called Wever-Bray effect. “Wever and Bray reported that if the electrical activity picked up from the cat's auditory nerve is amplified and redirected to a loudspeaker, then one can talk into the animal's ear and simultaneously hear himself over the speaker.” Gelfand (1998), p. 136.

frequency change. Consider two types of psychoacoustic experiment.⁹ The first type consists in ‘*Fractionalization*’ experiments. Subjects are presented with a given tone, and are then instructed to adjust a second (simultaneously presented) tone so that its pitch is one-half that of the first tone.¹⁰ This process is repeated for many tones of different frequencies. The second type are called ‘*Equal Intervals*’ experiments. In one such experiment, subjects are instructed to adjust the frequencies of five tones until they are separated by equal pitch intervals.¹¹ Fractionalization and Equal Intervals experiments yield a remarkably consistent scaling of pitch as a function of frequency, according to which equal pitch intervals do not correspond to equal frequency intervals.¹² Doubling frequency does not uniformly double pitch: the frequency of a 1000 Hz tone must be tripled in order to double its pitch; doubling the pitch of a 2000 Hz tone requires quadrupling its frequency. The accepted pitch scale extracted from psychoacoustic data assigns to equal pitch intervals equal magnitudes in units of *mels*. The mel scale is therefore an *extensive* or *numerical* pitch scale, in contrast to the *intensive* frequency scale for pitch.¹³ (See Figure 3: Pitch in *Mels* as a Function of Frequency, p. 28).¹⁴

Suppose we accept, with most auditory researchers, that the notion of pitch magnitude is well-founded, and that Fractionalization and Equal Intervals experiments reliably determine the relationship between perceived pitch and frequency.¹⁵ If pitch perception is for the most part veridical, it follows

⁹As described in Gelfand (1998), p. 354.

¹⁰Stevens, Volkman, and Newman (1937).

¹¹Stevens and Volkman (1940).

¹²Stevens, et al. (1937) and Stevens and Volkman (1940).

¹³Extensive scales preserve ratios between quantities, but intensive scales need not.

¹⁴Another accepted pitch scale is slightly different. I shall discuss, beginning in §5, the *Bark* scale introduced by Zwicker (1961) and Zwicker and Terhardt (1980).

¹⁵Yost and Watson (1987), Bregman (1990), Gelfand (1998), Mestre, et al. (1998), and Zwicker and Fastl (1999) are prominent recent examples of this view. Laming (1997), however, has argued against the assumptions that ground such psychophysical views. In particular, Laming has criticized the view that sensations themselves can be measured and assigned magnitudes:

The traditional idea of sensation is succinctly expressed by Luce (1981, p. 197): ‘... a physical stimulus is transduced into some sort of distinctive neural

that pitch is not identical with frequency. It also follows, perhaps surprisingly, that the octave relation is not a doubling of pitch. The octave relation is in some sense “the same again (but higher/lower)”, but it is not “double in pitch”. Likewise, none of the musical relations is a simple, small whole-number ratio between pitch magnitudes. We require another characterization of the octave, fifth, et al., as relations between sounds with particular pitches.

The Simple View is inadequate if we wish to identify pitch with some physical property of sounds that we veridically perceive. The relational structure among pitches differs from the relational structure of frequencies. In particular, changes in pitch magnitude do not correspond to uniform changes in frequency magnitude.

4 The Standard (Subjectivist) View

The accepted view among auditory researchers is that pitch is a *subjective* or *psychological* quality merely correlated with the frequency of a sound. Gelfand (1998), for instance, states that “in formal terms, pitch is the psychological correlate of frequency, such that high frequency tones are heard as being ‘high’ in pitch and low frequency tones are associated with ‘low’

activity which, under processing by the central nervous system, is ultimately perceived as a sensation.’ That is the idea I wish to abandon. I emphasize, instead, that it is the stimulus that is perceived, not the neural activity; and the stimulus is perceived as an object ‘out there’, not as an internal sensation (internal stimuli such as pain and tickle excepted). . . . The point I make is that the evidence so far to hand does not support any intermediate continuum at the psychological level of description which might reasonably be labelled ‘sensation’ (p. 205).

To this extent, Laming’s view agrees with one main lesson of this paper: that pitch is a property of sounds themselves, and not of sensations. I attempt, however, to provide an account of pitch that captures and predicts subjects’ patterns of judgment about what they have heard, on the assumption that those judgments reflect experiences that vary with a sound’s frequency as expressed in Figure 3. For this reason my view is strictly speaking opposed to Laming’s view that a “*natural* physical measure of stimulus magnitude” can fully capture sensory judgments (p. 167, my italics). According to the view I develop, to capture sensory judgments of pitch requires enlisting a more complex physical measure of stimulus magnitude.

pitches.”¹⁶ The *Standard View* among psychoacousticians is that pitch is a mental quality in virtue of which we imperfectly perceive the frequencies of sounds. Thus stated, the accepted view is that pitch is a quality of experiences much like a “quale” or “subjective feel” associated with frequency perception. The Standard View is a form of philosophical *subjectivism* that arises in response to the divergence between pitches and the external physical properties thought to be responsible for pitch experiences. Pitches are thought to be internal, mental properties, and thus subjective in a natural sense.

On the Standard View, sounds have frequency but not pitch. As such, it is a form of *error theory* concerning auditory perceptual experience. Since we take auditory experiences to furnish awareness of sounds and their pitches, if the Standard View is correct then pitch is what Alex Byrne has aptly described as “a perfectly monstrous illusion.”¹⁷ Our ascriptions of pitch properties to sounds are simply never true.

But it just does not follow from the fact that perceived pitch intervals do not correspond to like frequency intervals that pitch is a property of experiences and cannot be ascribed truly to sounds. I wish to propose that an *Alternative View* of pitch, according to which pitches are physical properties of sounds, is equally viable and captures more of the desiderata that should be met by a philosophical theory of sounds. If sounds have a physical property whose variations correspond uniformly to variations in perceived pitch and that is responsible for pitch experiences, then we can avoid both eliminativism about pitch and error theories concerning pitch perception.¹⁸

¹⁶Gelfand (1998), p. 353. Similarly, Mestre, et al. (1998) say, “pitch is the subjective quality associated with frequency.”

¹⁷Byrne uses this phrase to refer to the status of colors according to error theorists (Byrne (forthcoming), p. 9). Boghossian and Velleman (1989) and (1991), and Hardin (1988) are contemporary proponents of color eliminativism.

¹⁸What about other options that neither attribute systematic error to pitch perception nor eliminate pitches from the world of sounds? What we would like is a view that ascribes pitches to sounds and entails that most of the time, when things are going well with hearing, we perceive sounds to have the pitches they actually have. One such view is that pitches are dispositions, construed objectively, to cause subjects to have pitch experiences.

For now, we can summarize the respects in which the Standard View falls short of a fully satisfactory account of the audible qualities, and thereby illustrate the motivations for an alternative.

(1) If we accept that pitch is the property of sounds that is *causally responsible* for our experiences as of sounds' having pitch, and we accept that sounds are external, objective (mind-independent) entities, we will prefer an account according to which instances of the audible qualities are (1a) entirely *external* to perceivers and (1b) *objective* qualities that are instantiated in the world independently of perceivers and their responses. (1a) weighs against the Standard View, and (1b) also counts against other forms of subjectivism

We must take care in formulating dispositionalism. If pitches are dispositions to produce auditory experiences with a certain phenomenal character (compare, with respect to color, Peacocke (1984)), we have no direct perceptual access to the pitches of sounds since we are immediately aware only that we have an experience with certain phenomenal pitch properties. We can then infer or otherwise come to believe that sounds have dispositions to cause in us particular phenomenal pitch properties, but we will never be auditorily aware that some sound has a particular pitch. Byrne (ms) offers a more developed version of this argument against what he calls '*reductive dispositionalism*'. This form of dispositionalism does not leave us much better off than the Standard View.

Better to say that pitches are dispositions to produce experiences as of a sound's having a certain pitch. Byrne (ms) calls this view '*nonreductive dispositionalism*'. But the nonreductive dispositionalist must accept that auditory experiences present *relational* properties of sounds—e.g., *being disposed to sound high-pitched to perceivers of kind k*—that we are strongly inclined to think are monadic, or that auditory experiences do not present sounds as having pitches, properly understood. Neither disjunct is attractive, but perhaps nonreductive dispositionalism, suitably motivated, is capable of capturing most of what we want from a theory of pitch.

I do not hope to settle the question of whether pitches, and sensible qualities more generally, are dispositions. My aim is to present a view of pitch according to which pitches are physical properties of sounds, and to defend it against objections that purport to demonstrate its inability to account for important features of pitch and pitch experiences. If sounds have a physical property that is responsible for pitch experiences, variations in which correspond to variations in perceived pitch, then we can avoid eliminativism about pitch and error theories of pitch perception. Given that the physical *bases* of dispositions are central to understanding what dispositions are and when they are correctly ascribed (cf. Martin (1994), Lewis (1997), and Fara (2001)), my project should be of interest to the dispositionalist who aligns herself against the Standard View according to which pitch is a psychological or subjective property.

that assign distal locations to pitch instances.¹⁹

(2) The Standard View entails that sounds themselves are not immediately characterizable in terms of pitch. If sounds can be correctly described and classified by their pitches without reference to their frequencies or mention of perceivers' responses, no widespread-error theory, eliminativism, or projectivism is correct for pitch.

(3) Knowledge about at least some of the properties of sounds seems to be possible solely on the basis of (perhaps suitably idealized) perceptual experiences. If such knowledge is not acquired by inference, quasi-inference, or some other indirect method, then the Standard View and other error theories about the processes involved in perception and belief formation about the audible qualities are false.

(4) Many philosophers maintain that the contents of perceptual experiences are exhausted by their representational or intentional contents—that experiences can be fully characterized by how they present the world as being. Intrinsic qualities of experiences, in particular, need not be mentioned in a full characterization of experiential contents.²⁰ The Standard View, however, essentially makes reference to the psychological property of pitch to

¹⁹Might sounds indeed bear pitch properties that are external to subjects though still in some sense subjective or mind-dependent? If so, the extreme conclusions of the Standard View might be avoided while maintaining the subjectivity of pitch. The most familiar form of subjectivism about sensible qualities is formulated in terms of dispositions to produce characteristic kinds of perceptual responses (e.g., Locke (1975) and McGinn (1983)). Pitch, then, is the subjective property sounds have in virtue of causing subjects to have pitch experiences. But nothing about a sound's being disposed to produce pitch experiences makes it the case that its pitch "depends" in some interesting sense upon subjects. Rosen (1994) presents convincing arguments that to claim of certain properties that they are "response dependent" is to mistake a distinction at the level of concepts or representations for a distinction at the level of properties or facts.

Perhaps pitches are basic, subject-dependent properties of sounds that cannot be identified with dispositions. But that makes pitches and their subjectivity, alike, mysterious from the perspective of a moderately naturalistic understanding of the world. The Standard View thus appears to be the most viable view that is in some way subjectivist and which meets the standards of explanation required by auditory science.

²⁰See, e.g., Harman (1990), Dretske (1995), Lycan (1996), Tye (1995) and (2000), and Byrne (2001).

account for the character of subjects' auditory experiences. If intentionalism is the correct *theory of perception*, an explanation of the content of auditory experience that adverts to a non-psychological pitch property must be given.

(5) Perception seems to be a relatively reliable mode of access to information about the world. We think that beliefs formed on the basis of our perceptual experiences are at least *prima facie* justified. To this extent we have reason to resist the claim that auditory experiences are in some important respect illusory. A view that vindicates pitch experience is preferable to one that attributes to hearing some substantial measure of error, either in that we distort frequency relations or in that our pitch experiences fail to discern any real property of sounds.

5 The Alternative (Physicalist) View

Physicalism about pitch amounts to the claim that there are properties of sounds themselves that can be described in the terminology of a physical theory and which correlate well with the pitch experiences of normal human hearers who are in circumstances that are favorable for hearing the pitches of sounds. We need only look for an alternative to physicalism if such a property is not in the offing or if further philosophical considerations prevent us from identifying pitch with any property expressible in the language of physics.²¹ The Standard Subjectivist View is motivated by the implicit assumption that the psychophysical results discussed in §3 show that no physical property of sounds can be identified with pitch since the most likely candidate—frequency—cannot. I wish to challenge this implicit assumption by characterizing such a physical property.

Proponents of the Standard Subjectivist account of pitch classify sensory response types according to the frequencies and intensities of the sounds that give rise to those responses. In fact, Zwicker and Terhardt (1980) express

²¹One such consideration may arise out of cases of spectral shift or inversion of sensible qualities with respect to external physical properties. A full discussion of auditory spectral shifts is contained in §A.3 of my doctoral dissertation, *Sounds* (Princeton University, 2002) (available online at <http://people.ucsc.edu/~cjo/papers/Sounds.pdf>).

the subjective pitch of a sound as a function of its (fundamental) frequency, and derive from this function a pitch scale in units called ‘Barks’ which, in its relationship to frequency, closely resembles that of the mel scale.²² (See Figure 4: Pitch in *Barks* as a Function of Frequency, p. 29). But frequency is, of course, a physical property of sounds. The physicalist about pitch has as her target precisely a property of sounds that varies with frequency just as the subjectivist’s psychological pitch does. I shall argue in the next section (§6) that extant accounts of pitch experience provide the materials to characterize just such a property. Accepted subjectivist accounts, including Zwicker and Terhardt’s, explain pitch experiences in terms of the activities of different portions of the auditory system that respond to the presence of energy within ranges of frequency (*critical bands*) which vary in magnitude throughout the audible spectrum. (See Table 1: Selected Critical Bands, p. 30). But such responsiveness on the part of the auditory perceptual system is responsiveness *to* the presence of energy within certain perceptually salient frequency ranges since the responses vary in proportion to critical band energies. The energy distribution across an ordering of those frequency ranges is thus a good candidate upon which to base a realist account of pitches as properties of sounds. Characteristics of the distribution of energy across ordered frequency ranges are perfectly objective properties that can be ascribed to sounds independently of considerations about perceivers.

The physical property of sounds that serves as pitch candidate is thus causally responsible for the activity of the auditory perceptual system that subserves pitch experience. Energy within a particular range of frequency causes activity proportional to that energy in a portion of the auditory system which is tuned to that frequency range. Since that kind of activity is the basis of pitch experience, pitch experiences are perceptual responses to the physical pitch candidate. But since particular pitch experiences are not simply responses to the frequencies of sounds, pitch perception is not just

²²The function from frequency to pitch in Barks (*b*) is:

$$b = 13 \arctan(0.76f/1000Hz) + 3.5 \arctan(f/7500Hz)^2.$$

frequency misperception. Pitch is therefore not the “psychological correlate” of frequency.

If the physicalist proposal is correct, then certain sounds have pitch in addition to frequency and intensity. Pitch is not identical to frequency, but according to the current proposal, the pitch of a sound is intimately related to its (fundamental) frequency, just as frequency is related to positions over time. Pitch is thus a property that depends upon the frequencies of a sound's simple components—given a full characterization of the attributes of the frequency ranges to which audition is sensitive, the pitch of a sound can be determined from its frequency constituents. Pitch, that is, can be computed as a complex function of frequency. Zwicker and Terhard's function thus specifies the relationship between the frequency and the objective pitch of a sound.

Pitch, however, is anthropocentric. Humans and creatures like us are sensitive to pitch in virtue of how our auditory systems are arranged, but pitch is not terribly interesting from either a physicist's point of view or from the perspective of giving the simplest complete characterization of sounds. It may be that pitch perception is more efficient from the standpoint of designing the mechanisms for auditory perception, or is more useful for our needs than frequency perception would be. Neither the fact that pitch is not useful in scientific discussions of sound, nor the fact that only creatures with a given type of sensory apparatus can detect pitch, however, implies that pitches are not real properties of sounds that can be captured in physical terms.

The account I shall develop, in terms of weighted energies within an ordering of frequency ranges, can be extended in several important directions. First, to give an account of the pitches of complex sounds; second, to give an account of the musical relations that hold between sounds in virtue of pitch; third, to explain how timbre and loudness might be objective properties of sounds. Though I restrict my attention in this paper to pitch, the physical properties upon which the account of pitch is based play a prominent role in demonstrating that the audible qualities need not be considered mere psychological correlates of the objective properties—frequency, spectral com-

position, intensity—in terms of which sounds are ordinarily characterized.²³

So, according to the alternative, pitch is an objective property of sounds that is not identical with (fundamental) frequency. Pitch is the physical property that disposes sounds to produce experiences as of pitch in suitably equipped perceivers. This is the property in virtue of which sounds that differ in timbre and loudness can be equivalent in “height”, and in virtue of which periodic or musical sounds can be ordered according to the ratio scale obtained by psychophysical methods. Particular sounds have pitch in addition to frequency, though pitch is the more salient property from the point of view of auditory perception. Pitch lacks the naturalness of frequency, and is thus interesting only from an anthropocentric perspective.

6 Pitch and Critical Bands

What, then, is pitch? Consider pitch, expressed in mels, as a function of frequency (Figure 3). The Standard (Subjectivist) View holds that this reflects the relationship between frequency as an objective property of sounds and pitch as an attribute of experiential states. But I shall claim that nothing about the case should keep us from thinking that both pitch and frequency are objective properties of sounds. The plausibility of this claim depends upon whether we can provide a candidate property that varies as pitch does with frequency and which explains why pitch is related to frequency as it is. The function from frequency to pitch would thus inform us about the relationship between two different properties of sounds, rather than about the relationship between properties of sounds and of sensations.

Recent accounts of pitch provide insight into both why the standard view is a subjectivist view, and into how an objectivist theory of pitch should be developed.²⁴ Zwicker and Terhardt (1980) have developed an instructive account of the relationship between frequency and pitch. They have derived a function from frequency to subjective pitch, expressed in units called ‘Barks’,

²³I extend the account of pitch to provide a novel account of the musical relations, and address timbre and loudness in “Audible Qualities,” and in Ch. 3 and App. A of *Sounds*.

²⁴The scientific details of discussion to follow are drawn primarily from Zwicker and Fastl (1999), especially Chapters 4 through 7, Gelfand (1998), and Zwicker and Terhardt (1974).

which strongly resembles that of the mel scale.²⁵ (See Figure 4: Pitch in *Barks* as a Function of Frequency, p. 29). The function from frequency to pitch in Barks (b) is:

$$b = 13 \arctan(0.76f/1000Hz) + 3.5 \arctan(f/7500Hz)^2.$$

The resulting Bark scale, like the mel scale, preserves perceived ratios and assigns the same magnitudes to pitch differences judged to be equivalent. Now, the Bark scale is based on the notion of a *critical frequency band*. Critical bands are supposed to be psychologically real entities that explain the perceived pitch of a tone given its frequency and energy properties. A critical band is characterized by a frequency range (its *critical bandwidth*) around a given *center frequency*, to which that band is responsive. Evidence for the existence of critical bands comes from experiments involving *masking*, where one tone is used to interfere with a subject's ability to perceive or detect the presence of another test tone. Masking experiments provide significant evidence for critical bands because the capacity one tone has to mask another is not independent of their respective frequencies. In particular, a masking tone whose frequency is near to that of a test tone more effectively masks the test tone than does a masking tone whose frequency is farther from that of the test tone. The effectiveness of masking decreases as frequency separation increases. This gives evidence of a sort of frequency selectivity within the auditory processing system. That is, as discussed in §2.2, different portions of the auditory system deal with particular frequency ranges. When tones are near in frequency, masking occurs because the sensation produced by one tone interferes with that of the other. As frequency separations increase, different portions of the auditory system carry information about the sounds, so less interference results and masking is less effective. Critical bands depend upon the tonotopic organization of the auditory system.

Critical bands are ordinarily thought of as like filters with a roughly triangular shape around a given center frequency. That is, they pass the most energy at the center or *characteristic* frequency, and pass less energy as

²⁵The term 'Bark' is derived from the name of Barkhausen, an auditory scientist who studied the relationship between loudness and intensity (Zwicker and Fastl (1999), p. 160).

a tone deviates from this center frequency. Beyond the critical bandwidth, they pass no energy. According to this model, the auditory system is like a bank of many overlapping filters that are responsive to different frequency ranges and whose outputs determine the qualities a sound is heard to have.

Critical bands and their characteristic properties can be discerned by a number of experimental procedures.²⁶ One such procedure discerns critical bandwidth by determining the minimum bandwidth of noise required to maximally mask a simple tone at the center frequency. For a test tone at a given center frequency, a narrow band of noise around the center frequency begins to mask the test tone. As the noise band widens, it further masks the test tone until a bandwidth is reached beyond which further noise does not contribute to masking. This is said to be the critical bandwidth. Now, the effectiveness with which noise or a masking tone interferes with the detection of a centered test tone drops off with departure from the center frequency until the critical bandwidth is exceeded and no further masking results. So, the critical bandwidth is that bandwidth of noise that masks a simple tone at a given frequency just as effectively as wideband white noise. Because frequency components farther from the center frequency contribute less to masking than those nearer to the center frequency, each frequency within a critical band can be assigned a weighting with respect to the relative importance to masking of energy at that frequency.

From the perspective of understanding pitch, critical bands are very significant for two reasons. First, critical bandwidth varies greatly with frequency. (See Figure 5: Critical Bandwidth as a Function of Frequency, p. 31). Critical bandwidth is approximately 100 Hz for frequencies up to approximately 500 Hz, but above 500 Hz a rough estimate of critical bandwidth is 0.2 times the center frequency. Each critical band deals with a unique range of frequencies, and that range increases substantially as center frequency increases.

This leads to the second significant (and surprising) result: critical bandwidth correlates very well with pitch. Critical frequency bands simply correspond to nearly equivalent pitch distances (approximately 100 mels per

²⁶See Zwicker and Fastl (1999), Ch. 6, for a survey of such procedures.

critical band or one Bark). Whether we take a critical band centered at 1000 Hz (critical bandwidth 160 Hz) or one centered at 5000 Hz (critical bandwidth 1000 Hz), critical bandwidths amount to equivalent pitch distances.²⁷

From these two facts, we can characterize the difference in pitch between two frequencies, $\Delta p(f_1, f_2)$, in terms of the function, $G(f)$, from center frequency to critical bandwidth.²⁸

$$\Delta p(f_1, f_2) = \int_{f_1}^{f_2} G(x) dx.$$

Pitch differences are a function of the critical bandwidths and center frequencies of the various experimentally determined critical bands.

But because the basilar membrane is thought to be an important part of the basis for both critical bands and pitch perception, and because the basilar membrane is organized such that less physical space is assigned to a given frequency range as frequency increases (see footnote 27), researchers posit that there are only a finite number of critical bands whose center frequencies become more widely spaced as frequency increases. That is, each unit of physical space along the basilar membrane subserves the same number of critical bands, but increasing frequency ranges, so the center frequencies of critical bands are separated by increasing frequency differences. Critical band center frequencies decrease in density along the frequency scale. The Bark scale thus assigns one unit to each critical bandwidth. Each unit of Barks corresponds to the same number of centered critical bands. (See Table 1: Critical Bands, p. 30).

What's significant, again, is that the procedure does not begin by determining the different frequency ranges that correspond to equal pitch ranges.

²⁷In addition, each Bark corresponds to rough 27 units of the minimum detectable pitch difference or *just noticeable difference* between tones; a unit of just noticeable frequency difference, however, grows substantially as frequency increases. Just noticeable frequency differences range from ~ 2 Hz at very low frequencies to nearly 200 Hz at high frequencies (Zwicker and Fastl (1999), p. 161). Critical bands also correspond to roughly equal distances along the basilar membrane or cochlear partition (1 Bark ≈ 1.3 mm), whereas equal cochlear distances correspond to increasing frequency ranges from apex to base (0.2 mm ≈ 15 –20 Hz at apex, 0.2 mm ≈ 500 Hz at base (Zwicker and Fastl (1999), p. 160)).

²⁸I owe thanks to Adam Elga for discussion of how to formulate this expression.

Rather, it begins by discerning the masking relationships between tones of different frequencies. These masking relationships show that critical bands exist and make apparent the frequency selectivity of the auditory system. But that critical bands discerned in this way should amount to equal pitch intervals is anything but evident. It is nothing short of surprising that an accurate pitch scale can be obtained simply by assigning one unit of critical-band rate to each experimentally discerned critical frequency band.

Zwicker and Terhardt's Bark scale is thus an expression of pitch in intervals that correspond to the psychoacoustically determined critical bands, and the result strongly resembles that of the mel scale of ratio pitch. The Bark scale yields an expression of subjective pitch as a function of frequency, but also explains this relationship in terms of the psychological and physiological mechanisms of sound perception.

Now, according to the "filters" conception of critical bands, the experienced attributes of a sound depend upon the output or *critical band level* (in terms of energy) of each of the critical bands in the filter bank. What is distinctive about the critical band levels caused by sounds that are heard to have pitch? How do they differ from those produced by sounds experienced to be unpitched?

A wideband white noise whose energy is independent of frequency results in a roughly equivalent critical band level for each of the overlapping critical bands. If the *critical band profile* for a sound is a plot of the critical band levels that sound produces for each critical band, the critical band profile of white noise is a flat line at some energy level. However, for a narrower band of noise, critical bands within which some portion of that energy falls have greater critical band levels than critical bands within which none of the noise falls. The resulting critical band profile is that of a plateau. But noise has no pitch. Sinusoids are the simplest pitched sounds, and are the constituents of more complex pitched sounds. *Sinusoids produce a distinctive critical band profile.* A sinusoidal tone at a particular frequency causes a substantially greater critical band level for critical bands centered near that frequency than for any other critical band. The critical band profile for a sinusoid is that of a sharp peak at the critical bands centered closest to its frequency. If

the experienced attributes of a sound depend upon its critical band profile, then the apparent pitch of a sinusoid depends upon its producing a maximum critical band level in a very small subset of adjacent critical bands among the many ordered critical bands.²⁹

A proponent of the Standard View might identify pitches with critical band maxima or with something that depends upon such activity. Either way, the pitch that a sound appears to have depends upon the profile of energies across many critical bands. Pitch is thus either identifiable with or the immediate result of a certain type of critical band activation pattern. Since critical bands are internal psychological entities, pitch is an internal subjective property.

So the Standard View posits internal psychological entities—critical bands—to explain, among other things,³⁰ the dependence of masking on frequency separation. A function from frequency to subjective pitch can be derived just by taking into account the experimentally determined attributes of the sequence of critical bands. The pitch of a sinusoidal tone depends upon the pattern of critical band activations it produces. The critical band profile associated with a sinusoidal tone is that of a peak critical band level in a single critical band or in a few adjacent critical bands. Thus, subjective or psychological pitch depends upon a critical band's having a significantly greater level of activity than the others. This critical band's location in the ordering of critical bands determines the Bark value associated with the tone.

But I claim that something analogous to critical band levels or profiles can also be ascribed to sounds themselves. If that is correct, then pitch can be seen as an objective property that sounds possess. We begin with the straightforward recognition that critical bands are themselves characterized in terms of a center frequency, a bandwidth, and a scaling factor for energy at each frequency within the critical bandwidth. But frequency and energy or

²⁹Indeed, very narrow band noise is often heard to have a pitch, in particular for noise bands at high frequency.

³⁰Critical bands also figure prominently in accounts of loudness. I discuss this fully in “Audible Qualities” and §A.4 of *Sounds*.

intensity are properties of sounds. So if we take a *critical band* to be a simple frequency range around a given center frequency, then sounds themselves fall within critical bands in virtue of their frequency characteristics. Given the scaling factor associated with each frequency within a critical band, a critical band can be characterized as a class of pairs that consist in a frequency and an energy weighting factor. The highest energy factor determines the center frequency of the critical band, and energy weighting factors decrease with distance from the center frequency. Now we can ascribe to every sound a *critical band profile* in the following manner. Ascribe to the sound a *critical band level* for each critical band: the critical band level of a sound, for a given critical band, is the sum of the products of the sound's energy at a frequency and the scaling factor for that frequency, for each frequency within the critical band. The critical band profile for a sound is given by ordering the critical band levels for all critical bands; it specifies the amount of energy that the sound has within each of the fully characterized critical bands.

Once we determine the critical band profile of a sound, we have a basis from which to determine its various attributes, including its pitch, loudness, and timbre. As before, for a sinusoidal tone to have a particular pitch is for it to have a single critical band maximum within its critical band profile. The critical band with the peak critical band level determines the sound's pitch in Barks.

Pitch is thus an objective though complex property of sounds. Having a pitch is not a simple matter of having a particular frequency, though frequency determines pitch given a specification of critical bands. The account depends upon rejecting the claim that pitches, and critical band profiles in general, are just properties of the auditory perceptual systems of subjects. Rather, the alternative holds that pitches and critical band levels are physical properties of sounds that our auditory apparatus is suited to detect. Activity internal to our perceptual systems reflects the presence of energy within various critical bands; that is, various parts of our auditory physiology and the psychological mechanisms they ground are tuned to energy within different frequency ranges. So the property of having energy within the specified critical bands is the property discerned in sound perception. Having a criti-

cal band level sufficiently greater than other critical band levels amounts to having a pitch.

An analogy with a device that determines the acceleration of an object illustrates the essence of the contrast between the Standard View and the alternative I have proposed. Acceleration is a property that an object has which can be calculated as a function of its position over time—acceleration depends upon positions and times, which look to be the simplest observable properties of an object that characterize its activity. Now, we are not inclined to say that the acceleration detector measures a property whose nature depends upon the construction of the device, in the sense that there could be no such property if the device did not exist. Once we recognize that some property of the object corresponds to the device's output, there is no temptation to say that the calculated acceleration is a property of the detecting device or that acceleration metaphysically depends upon detecting devices. Acceleration is a property of the object which the device is designed to detect. Similarly, once we recognize that critical band properties of sounds correspond to pitch experiences, there is little temptation to identify pitch as the mere subjective correlate of frequency. Pitch experiences track a property of sounds that varies with frequency.

Pitch is thus a complex property of sounds that depends upon their patterns of critical band levels. It is given by a complex function of the frequencies at which a sound has energy, and the function from frequency to pitch in Barks expresses the pitch of a sound itself—not merely a sensation. Since this function is derived from the attributes of the critical bands, and because critical band properties can be ascribed naturally to sounds, pitch is a property of sounds that is related to but not identical with frequency.

6.1 The Pitch of Complex Tones

The foregoing account of pitch in terms of critical band energies applies straightforwardly to simple sinusoidal tones, for which a single critical band has a maximum value. But complex sounds with many sinusoidal constituents also have pitch. The standard account of pitch perception for complex tones has it that fundamental frequency is determined through analysis

of a tone's Fourier frequency components. How can the alternative account I have sketched explain pitch perception for complex tones without adverting to fundamental frequency perception?

Recent accounts of pitch perception posit a complicated analysis of constituent frequencies to extract fundamental frequency. I shall consider one such model of pitch perception—that of Terhardt, e.g., (1980), (1982a), and (1982b)—and argue that it can be extended to the critical bands account of pitch developed above. The result is that pitch perception for complex sounds involves determining a critical band value in Barks; but the critical band value itself corresponds roughly to fundamental frequency.

Complex sounds have more widely distributed critical band energy patterns than do simple sinusoids, but complex periodic tones still exhibit local maxima at several critical bands. Pitches are therefore best conceived as types of critical band profiles. A sound whose critical band profile has many separate peaks corresponding to its sinusoidal constituents can be a member of the same pitch as one whose critical band profile has a single peak. The first problem is how to determine which complex sounds are members of the same pitches as sounds that have single-peaked critical band profiles. A simple answer is that they share a fundamental frequency, but we would like an answer framed in terms of aspects of critical band profiles. The second problem is how, if we are not perceiving fundamental frequency (because we do not perceive frequency at all), we come to perceive simple and complex sounds as having the same pitch. I shall suggest that Terhardt's account provides the materials to deal with both problems.

Call each local maximum a *pitch determinant* with a particular Bark value. Each pitch determinant corresponds to a particular frequency; for harmonic complex sounds—ones with pitch—these frequencies are integer multiples of each other. No such simple relationship exists, however, between the corresponding Bark values of pitch determinants. For instance, a complex harmonic sound with 200 Hz fundamental frequency and harmonics at 400 Hz, 600 Hz, and 800 Hz (or f , $2f$, $3f$, $4f$) has pitch determinants with Bark values of roughly 2 Barks, 4 Barks, 5.75 Barks, and 7.2 Barks. Terhardt's account of pitch perception relies on learned “templates” for particular har-

monic sequences, which are formed on the basis of regular experience with harmonic complexes, in particular those which are predominant in speech.³¹ We form the ability to recognize particular harmonic complexes, such as (150 Hz, 300 Hz, 450 Hz, 600 Hz, ...), through their spectral constituents. We thereby acquire internal templates with such constituents as components. But each component of a template has a corresponding value in Barks, so the templates can themselves be considered as sequences of pitches in Barks (e.g., (B1, B2, B3, B4, ..., B_n)) as opposed to frequencies (with the form (f, 2f, 3f, 4f, ..., nf)). Templates construed in such terms amount to sets of critical band maxima.

Each constituent of a complex tone is a member of various templates in which it occurs as an upper harmonic. Terhardt's claim is that an analysis on the subharmonics of each constituent of a complex tone determines the pitch that tone is perceived to have.³² The analysis proceeds as follows. For each of the n constituents of a complex tone, the first eight subharmonics of f_n are given by the fundamentals of the first eight harmonic sequences or templates of which f_n is an upper harmonic.³³ Consider the complex sound with constituents at 200 Hz, 400 Hz, 600 Hz, and 800 Hz. The subharmonics for 200 Hz are first determined: 100 Hz is the second subharmonic (1:2) of 200 Hz since 200 Hz is the second upper harmonic in the sequence (100 Hz, 200 Hz, 300 Hz, ...); 66.7 Hz is the third subharmonic (1:3) since 200 Hz is the third upper harmonic in the sequence (66.7 Hz, 133.3 Hz, 200 Hz, 266.7 Hz, ...). The final five subharmonics of 200 Hz are given by the ratios 1:4–1:8. Subharmonics are then determined for each of the complex tone's constituents.

A complex periodic sound's constituents share various subharmonics. Call a shared subharmonic a *pitch candidate*. Though several subharmonics

³¹Such learning is thought to take place very early in life, and probably begins in the womb. See *Sounds*, §A.2 for discussion of the basis of template formation—of why such templates are acquired.

³²I will adopt, for simplicity, the non-standard convention of including a tone itself among its subharmonics and upper harmonics.

³³So, for f_n , the first eight subharmonics are given by g_m where $f_n = 1g_1$, $f_n = 2g_2$, $f_n = 3g_3$, ..., $f_n = 8g_8$.

may coincide for two or more of the constituents, and therefore determine several pitch candidates, the pitch candidate shared by the largest number of constituents determines the perceived pitch of a complex sound. A comparison across subharmonics reveals that the fundamental frequency is shared as a subharmonic by the most constituents. That is, 200 Hz, 400 Hz, 600 Hz, and 800 Hz all share only 200 Hz as a subharmonic. Thus, 200 Hz determines the perceived pitch of the complex sound. Sometimes other pitch candidates are selected as “the pitch” of a complex sound, and often other pitch candidates are discernible within the complex sound. However, a complex sound is most commonly pitch-matched with a sinusoid at its fundamental frequency. This analysis explains both the “spectral” pitch of complex sounds with a constituent at the fundamental frequency, and the “virtual” pitch of sounds with a missing fundamental.

I claim that this analysis can be carried out with templates consisting not of frequencies in whole-number multiple relationships, but of pitch values in Barks which are derived from the critical bands model. Suppose we have a complex tone consisting of constituents at 2 Barks, 4 Barks, 5.75 Barks, and 7.2 Barks. Each of these constituents appears in learned templates derived from experience with complex harmonic sounds. Just as with the frequency model, coincidence among the subharmonics of the pitch determinants gives the pitch candidates and the pitch of the complex sound (which corresponds roughly to the fundamental frequency). Since templates are just representations of sequences of simple pitches, the pitch of a complex sound can be seen as settled by the pitches of its constituents and their subharmonics, which are given by simple harmonic sequences of *critical band values*. Having a particular pitch, for a complex sound, is a matter of having critical band maxima that substantially coincide in their membership in other salient (harmonic) sequences of simple pitches (alternatively, salient critical band profiles). Thus, though the pitch value of a complex sound has a corresponding frequency value that roughly equals the sound's fundamental frequency, what we perceive is its pitch. The point of the subharmonics analysis is to determine which harmonic sequence each of the complex sound's constituents are most likely to be members of. Pitches, then, can be seen

as types of critical band profiles whose members exhibit a particular kind of subharmonic sharing. The subharmonic shared determines the pitch value in Barks. Though pitch perception for complex sounds may be a quite complicated affair, it need not rely on determining fundamental frequency.³⁴

7 Concluding Remarks

I have argued that the failure of simple forms of physicalism about the audible qualities does not rule out the prospect of an account according to which pitch is an objective, physical property of sounds. The motivation for construing pitch as a subjective or psychological quality stemmed from the discrepancy between frequency and perceived pitch. But with aid of current accounts of pitch experience, we can direct our attention toward the physical property that both causes pitch experiences and varies with perceived pitch in the right ways. Explanation of pitch perception need not advert to mind-dependent properties: pitch is a property of sounds—which sounds have in addition to frequency—that is discerned during auditory experiences. Pitches as the alternative characterizes them are not, however, terribly interesting from the point of view of the physical sciences: frequency is a more natural property by which to classify sounds from the perspective of the physical sciences. But pitches are salient from our perspective as hearers, and add intelligibility to the complex mechanisms by which we hear sounds. Pitch is thus an anthropocentric property.³⁵ But the property of having a pitch—that is, having or determining a maximal energy within a given member of an ordering of frequency ranges—is neither mental nor mind-dependent in any important respect.³⁶

³⁴A further benefit of this account of pitch is that it provides an analysis of the musical relations. Patterns of subharmonic sharing—a prominent feature in the account of the pitch of complex sounds—form the basis of the musical relations. In *Sounds*, §A.2, I develop this account.

³⁵Just as colors are anthropocentric properties according to the view that colors are types of surface spectral reflectances. See, e.g., most recently, Bradley and Tye (2001), and Byrne and Hilbert (forthcoming in *BBS*).

³⁶Thanks, in particular, to Adam Elga, Gilbert Harman, Gideon Rosen, and Jeffrey Speaks for discussion and comments on the material in this paper.




Figure 1: Sinusoidal motion. [From Gelfand (1998), *Hearing: An Introduction to Psychological and Physiological Acoustics*, Third Edition, New York: Marcel Dekker, Figure 1.5, p. 15, with permission.]




Figure 2: Fourier composition of a square wave. [From Gelfand (1998), *Hearing: An Introduction to Psychological and Physiological Acoustics*, Third Edition, New York: Marcel Dekker, Figure 1.12, p. 24, with permission.]




Figure 3: Pitch in *mels* as a function of frequency. [From Gelfand (1998), *Hearing: An Introduction to Psychological and Physiological Acoustics*, Third Edition, New York: Marcel Dekker, Figure 12.1, p. 354, with permission.]




Figure 4: Pitch in *Barks* as a function of frequency. [From Gelfand (1998), *Hearing: An Introduction to Psychological and Physiological Acoustics*, Third Edition, New York: Marcel Dekker, Figure 12.2, p. 355, with permission.]

b Bark	f_l, f_u Hz	f_c Hz	b Bark	Δf_G Hz	b Bark	f_l, f_u Hz	f_c Hz	b Bark	Δf_G Hz
0	0				12	1720			
		50	0.5	100			1850	12.5	280
1	100				13	2000			
		150	1.5	100			2150	13.5	320
2	200				14	2320			
		250	2.5	100			2500	14.5	380
3	300				15	2700			
		350	3.5	100			2900	15.5	450
4	400				16	3150			
		450	4.5	110			3400	16.5	550
5	510				17	3700			
		570	5.5	120			4000	17.5	700
6	630				18	4400			
		700	6.5	140			4800	18.5	900
7	770				19	5300			
		840	7.5	150			5800	19.5	1100
8	920				20	6400			
		1000	8.5	160			7000	20.5	1300
9	1080				21	7700			
		1170	9.5	190			8500	21.5	1800
10	1270				22	9500			
		1370	10.5	210			10500	22.5	2500
11	1480				23	12000			
		1600	11.5	240			13500	23.5	3500
12	1720				24	15500			
		1850	12.5	280					

Table 1: Selected critical bands. Bark value b , lower (f_l) and upper (f_u) frequency limit of critical bandwidths, Δf_G , centered at f_c . [Adapted from Zwicker and Fastl (1999), *Psychoacoustics: Facts and Models*, Second Edition, New York: Springer-Verlag, Table 6.1, p. 159, with permission.]




Figure 5: Critical bandwidth as a function of frequency. [From Zwicker and Fastl (1999), *Psychoacoustics: Facts and Models*, Second Edition, New York: Springer-Verlag, Figure 6.8, p. 158, with permission.]

References

- [1] Boghossian, Paul A. and Velleman, J. David. (1989). "Colour as a Secondary Quality." *Mind* 98: 81–103.
- [2] ——. (1991). "Physicalist Theories of Color." *Philosophical Review* 100 (1991): 67–106.
- [3] Bradley, Peter and Tye, Michael. (2001). "Of Colors, Kestrels, Caterpillars, and Leaves." *Journal of Philosophy* 98: 469–487.
- [4] Bregman, Albert S. (1990). *Auditory Scene Analysis: The Perceptual Organization of Sound*. Cambridge, MA: MIT Press.
- [5] Byrne, Alex. (2001). "Intentionalism Defended." *Philosophical Review* 110: 199–240.
- [6] ——. (forthcoming). "Color and Similarity." *Philosophy and Phenomenological Research*.
- [7] ——. (ms). "Colors and Dispositions." Unpublished manuscript.
- [8] Byrne, Alex, and Hilbert, David (eds.). (1997). *Readings on Color, Volume 1: The Philosophy of Color*. Cambridge, MA: MIT Press.
- [9] ——. (forthcoming BBS). "Color Realism and Color Science." *Behavioural and Brain Sciences*.
- [10] Dretske, Fred. (1995). *Naturalizing the Mind*. Cambridge, MA: MIT Press.
- [11] Fara, Michael. (2001). *Dispositions and Their Ascriptions*. Princeton University Doctoral Thesis.
- [12] Gelfand, Stanley A. (1998). *Hearing: An Introduction to Psychological and Physiological Acoustics*. 3rd ed. New York: Marcel Dekker.

- [13] Hardin, C.L. (1988). *Color for Philosophers: Unweaving the Rainbow*. Indianapolis: Hackett.
- [14] Harman, Gilbert. (1990). "The Intrinsic Quality of Experience." In James Tomberlin (ed.), *Philosophical Perspectives* 4. Atascadero, CA: Ridgeview.
- [15] Helmholtz, Hermann. (1954). *On the Sensations of Tone*. New York: Dover. (Translated conforming to the Fourth German Edition of 1877 by Alexander J. Ellis).
- [16] Laming, Donald. (1997). *The Measurement of Sensation*. New York: OUP.
- [17] Lewis, David. (1997). "Finkish Dispositions." *The Philosophical Quarterly* 47: 143–58.
- [18] Locke, John. (1975). *An Essay Concerning Human Understanding*. Peter Nidditch (ed.). Oxford: Clarendon Press.
- [19] Luce, R.D. (1981). "Physical Correlate Theory: A Question and a Prediction." *Behavioral and Brain Sciences* 4: 197–198.
- [20] Lycan, William. (1996). *Consciousness and Experience*. Cambridge, MA: MIT Press.
- [21] Martin, C.B. (1994). "Dispositions and Conditionals." *The Philosophical Quarterly* 44: 1–8.
- [22] McGinn, Colin. (1983). *The Subjective View: Secondary Qualities and Indexical Thoughts*. Oxford: Clarendon.
- [23] Mestre, J., Mullin, W.J., and Gerace, W.J. (1998). "Course Notes for Physics 114, Theory of Sound with Applications to Speech and Hearing Science." <http://www-unix.oit.umass.edu/~phys114>.
- [24] Peacocke, Christopher. (1984). "Colour Concepts and Colour Experience." *Synthese* 58: 365–382.

- [25] Rosen, Gideon. (1994). “Objectivity and Modern Idealism: What is the Question?” In Michaelis Michael and John O’Leary-Hawthorne (eds.), *Philosophy in Mind*. Dordrecht: Kluwer.
- [26] Schouten, J.F. (1940). “The Residue, A New Concept in Subjective Sound Analysis.” *Proceedings of the Koninklijke Nederlandse Akadademie* 43: 356–365.
- [27] Stevens, S.S. and Volkmann, J. (1940). “The Relation of Pitch to Frequency: A Revised Scale.” *American Journal of Psych.* 53: 329–353.
- [28] Stevens, S.S., Volkmann, J., and Newman, E.B. (1937). “A Scale for the Measurement of the Psychological Magnitude Pitch.” *Journal of the Acoustical Society of America* 8: 185–190.
- [29] Terhardt, Ernst. (1974). “Pitch, Consonance and Harmony.” *Journal of the Acoustical Society of America* 55: 1061–1069.
- [30] ——. (1979). “Calculating Virtual Pitch.” *Hearing Research* 1: 155–182
- [31] ——. (1980). “Toward an Understanding of Pitch Perception: Problems, Concepts, and Solutions.” In G. van den Brink and F.A. Bilsen (eds.), *Psychophysical, Physiological, and Behavioral Studies in Hearing*. Delft: University Press, 353–360.
- [32] Terhardt, E., Stoll, G., and Seewann, M. (1982a). “Pitch of Complex Signals According to Virtual Pitch Theory: Tests, Examples, and Predictions.” *Journal of the Acoustical Society of America* 71: 671–678.
- [33] ——. (1982b). “Algorithm for Extraction of Pitch and Pitch Salience from Complex Tonal Signals.” *Journal of the Acoustical Society of America* 71: 679–688.

- [34] Tye, Michael. (1995). *Ten Problems of Consciousness*. Cambridge, MA: MIT Press.
- [35] ——. (2000). *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- [36] Woolsey, C.N. (1960). “Organization of the Cortical Auditory System: A Review and Synthesis.” In G.L. Rasmussen and W.F. Windle (eds.), *Neural Mechanisms of the Auditory Vestibular Systems*. Illinois: Thomas, Springfield, 165–180.
- [37] Yost, William A. and Watson, Charles S. (eds). (1987). *Auditory Processing of Complex Sounds*. New Jersey: Erlbaum.
- [38] Zwicker, Eberhard. (1961). “Subdivision of the Audible Frequency Range into Critical Bands (Frequenzgruppen).” *Journal of the Acoustical Society of America* 33: 248.
- [39] Zwicker, Eberhard and Fastl, Hugo. (1999) *Psychoacoustics: Facts and Models*, Second Updated Edition. New York: Springer-Verlag.
- [40] Zwicker, Eberhard and Terhardt, Ernst (eds.). (1974). *Facts and Models in Hearing*. New York: Springer-Verlag.
- [41] ——. (1980). “Analytical Expressions for Critical-band Rate and Critical Bandwidth as a Function of Frequency.” *Journal of the Acoustical Society of America* 68: 1523–1525.