Feeling a beat

Alex Kerr

A few years ago, you probably saw a picture of a dress. The dress looks blue to some and white to others. Since it can't be both colors, it can't be how it looks to everyone—the appearances conflict. A predictably heated debate followed about who sees things as they are.

You may have missed a similar flamewar about a song: "Videotape", by Radiohead. This debate was about the *beat* of the song: to a first approximation, an experienced pulse that we naturally tap to. To illustrate, please listen to the first twenty seconds.¹ Most feel a beat under the song's opening chords. But some feel a beat between the chords. The debate was about who hears things as they are—about where the beat is.

I'll soon help you hear "Videotape" both ways. But let's start with something simpler: a metronome.² Set the metronome's tempo to 100 beats per minute (bpm) and accent to none. Now hit play. You'll hear beeps. And, by default, you'll probably feel a beat under each sound. If you don't, try tapping along—this usually helps.

Now try to feel a beat between the beeps. If you can't right away, here is some training. First set your metronome's tempo to 200 bpm, its accent to two, and hit play. High and low beeps alternate. Tap on the low ones. You should start to feel a beat under the low beeps, between the high ones. Now your training is done. Replay the high beeps unadorned—100 bpm, no accent—and try to feel a beat between the high beeps, as before.

If all has gone well, you felt a beat under each sound, then between them. And, feeling a beat each way, things appear differently to you—what it is like for you to hear the sounds is different. "Videotape" appears differently to different people in much the same way: some feel a beat under chords, while others feel a beat between them.

¹Radiohead, "Videotape," released 2007, on *In Rainbows*, self-released, accessed May 24, 2023, https://youtu.be /WF83_PR2EsA.

²Metronome available here: https://www.8notes.com/metronome/

This is a contrast in how things appear in perception, or what it is like to perceive. And philosophers usually offer just a few diagnoses of such contrasts. The beat is puzzling because it raises problems for all of them.

The first diagnosis is that the appearances *conflict*: just as the dress can be blue or white, but not both, the beat can land under sounds or between sounds, but not both. So, when some feel a beat under the chords, and others feel a beat between the chords, things can't be how they appear to everyone—hence the parallel flamewars about the dress and "Videotape".

This diagnosis explains the internet comments. But it makes bad predictions about how things appear. If you see the dress as blue then white, the thing seems to change in color. But no matter how you feel a beat under the metronome, the same kinds of sounds seem to ring out at the same times—the metronome doesn't seem to change its tune. Why, unlike the dress, do the sounds appear differently without appearing to change? I'll argue that, if appearances conflict, there is no good answer.

The second diagnosis says the contrast is a *change of topic*: we each hear different but compatible features of the song. Compare: the dress looks close to you, held at arm's length, and far away to me, seen from across the room. Different features appear to us, but the appearances don't conflict: the dress can be close to you and far from me, just as it appears to us.

This diagnosis explains why the metronome doesn't seem to change its tune. But it makes other facts about feeling a beat mysterious. You can see the dress' distance from you and its distance from me at a single glance. But you cannot feel both beats under "Videotape" at once—you can experience only one at a time. Why, if they are simply different compatible features? Here again, there is no good explanation.

The third diagnosis is that our different ways of hearing "Videotape" are *changes in mental paint*. On this view, special features of our experiences, distinct from features we perceive or even seem to perceive, explain the difference in how things appear to us. It is controversial whether there is any mental paint. But canonical examples include the ways bodily sensations feel. The problem for this view is that feeling a beat affects how sounds appear to you. It doesn't seem like an unrelated accompaniment. Why, if it is having a feeling-like mental state while also hearing sounds? Again, there is no good answer.

So, the beat resists all three diagnoses. But on the usual views of perception, these are the only

options. For on *perceptual views*, things appear as they do only because of which objective, or mind-independent, things, events, and features we experience. And, on the usual alternative, the *mental paint view*, things appear as they do because of mental paint. The beat is accordingly a problem for the usual views of perception. It motivates a better view that allows for different ways of experiencing the mind-independent world that are not explained in terms of mental paint. My hope is that developing such a view in this specific context illustrates the benefits of the view more generally.

That is my general aim. But more specifically, I will propose that feeling a beat is experiencing events relative to *units* of time: periods of time and timepoints that divide the periods. These units are not qualities of experiences. But they aren't mind-independent things, events, or features either—they aren't, for example, properties of or relations between *the sounds* you hear when you feel a beat. For just as you can represent the same goings-on using one calendar or another, you can represent the same things, events, and features using different beats—the choice of unit doesn't distinguish ways the mind-independent world seems to be.

Here is my plan. In section 1, I'll give a targeted introduction to the beat using musical examples. Then, in sections 2-4, I'll argue that contrasts in how we feel a beat are not conflicting appearances, changes in topic, or changes in mental paint. In section 5, I will develop my own view, address challenges, and discuss implications for theories of perception.

1 Background

My topic is a contrast in feeling a beat. I just gave an example. But now I want to single out the beat in the example more carefully. Here it will help to distinguish some components of the beat and distinguish the beat from nearby topics.

First, let's distinguish beat from rhythm. You heard a pattern of durations between beeps, or their *rhythm*.³ But you also felt a beat. And this takes more. For you heard the same rhythm—the same pattern of durations between sounds—no matter whether you felt a beat under or between the sounds.⁴

³In standard usage, rhythms also include more: for example, patterns in the timing between abrupt or gradual changes in pitch, timbre, loudness, or other perceptible changes in sound.

⁴A more dramatic illustration of the distinction is that you can hear rhythm but fail to feel a beat entirely. *Dysrhyth*-

Still, beat and rhythm share some features. Like the rhythmic beeps, the beat you feel is a series of pulses. And, again like the beeps, these pulses repeat regularly, or *periodically*. So, as with any periodic series, we can characterize a beat by its *phases*, or states, and its *period*, or the duration between phases. Let's call such a series a *beat* and each pulse in the series an *individual beat* to disambiguate, when needed. And let's call the rate of pulses per unit time (usually measured in *beats per minute* or *bpm*) the *tempo* of the beat.⁵

So far, we've looked at the beat listeners naturally tap to. This is called the *tactus*. But, whenever we experience a tactus, we're very good at tapping at some related rates, but not others. We can tap, for example, at double or half the tempo of the tactus, but not at 1/7 the tempo. The usual explanation is that the listener experiences or at least somehow keeps track of other beats at the tempo she can tap to, but fails to experience or keep track of those she can't tap to. Such beats that divide a tactus are *subdivisions*.

We're now in a position to say what changes, and what doesn't, when you feel a beat differently under the metronome. At first, you experienced individual beats of the tactus coincide with each beep of the metronome. Let's call this experience *Downbeats*. Next, you felt a beat with this same tempo but whose phases aligned differently with the phases of the metronome. Here individual beats of the tactus coincided with pauses between beeps. And subdivisions dividing the tactus in half coincided with each beep. Let's call this experience *Upbeats*.

This contrast is in apparent *syncopation*: it is a change in whether you hear sounds as landing on the beat—as *unsyncopated*—or instead hear sounds as landing between beats, on subdivisions—as *syncopated*. The two ways of hearing "Videotape" are likewise a change in apparent syncopation. But before I return to the contrast, there is one more aspect of the beat we need to review.

So far, I've emphasized similarities between beat and rhythm. But a beat has more structure. It also has a *meter*. To illustrate, count by fours—"one, two, three, four"—at a steady pace, and repeat. Repeat enough, and you will feel a beat. But you will also naturally experience the beat "going by fours": in groups, called *measures*, that are four individual beats long. Each of these measures is an instance of a *four-beat* meter. Now count by threes, at the same pace: "One, two,

mics have this problem chronically. See Jessica Phillips-Silver et al., "Born to Dance but Beat Deaf: A New Form of Congenital Amusia", *Neuropsychologia*, 5 (April 2011): 961–69.

⁵I'll stick to evenly timed, or *isochronous*, beats. This is a big restriction: the beat is often fitted in fine-tuned ways to the expressive timing of a performance. But the simplification shouldn't matter for our purposes.

three". You will feel a beat at the same tempo. But here you naturally experience the beat "going by threes": in measures that are three individual beats long. Each of these measures is an instance of a *three-beat* meter.

Each chant prompts an experience of meter. But it also assigns *coordinates* to the meter. We can count beats per measure with chanted numerals. And beats under a "one" are *beat 1* of the measure, those under a "two" are *beat 2*, and so on.⁶ Finally, to represent subdivisions, we can expand the chants. Add an "and" between numerals: "one-and, two-and, three-and, four-and". The beats under "and"s are subdivisions. The subdivision after beat *one* is the "*and*" of one, the one after beat *two* the "*and*" of *two*, and so on.

Now, finally, back to "Videotape". Please listen to a live performance using the link in the note.⁷ At the start, most listeners feel a beat under each chord, grouped into four-beat measures. Then, new cues creep in, and listeners tend to hear the chords as pounding on subdivisions—they sound syncopated. After the change, listeners experience a tactus twice as fast as the old one. And they experience chords land on the "and" of beats 2 and 4—the subdivisions just after beats 2 and 4. So, for instance, the chord you initially heard starting a measure is now heard on the "and" of the fourth beat, at the end of a measure.

On the album version, these cues are subtler. So, most default to hearing the figure as unsyncopated throughout the song. But once you've heard the live version, you can also hear the figure as syncopated.

So, you can hear these sounds in two ways: as syncopated or unsyncopated. Heard each way, things appear differently: what it is like to hear the sounds is different. The question is: why do things appear differently—what explains this contrast in what it is like for you to hear the sounds?

I think this is a question for theories of perception. So I'll assess theories by how well they answer it and then offer my own. But why think this contrast is a change in how things appear in *perception*—in this case, in what it is like to hear the sounds? Skepticism here is reasonable—

⁶Counting is useful for thinking and talking about meter. But experiencing meter isn't counting. Consider the Neil Young songs "Harvest" and "Only Love Can Break Your Heart". Normal listeners hear the latter in a three-beat-meter and the former in a four-beat meter. But they need not count. And counting isn't enough. You can mutter "one, two, three", repeating, along to "Harvest". But if you're a normal listener, you're chanting faithlessly—you're stuck hearing "Harvest" in a four-beat meter.

⁷Radiohead, "Videotape," Great Stage Park, Manchester, June 2006. Recording here: https://youtu.be/1hgXZAH onLs?t=1966

especially since I'll draw broad conclusions from my case study.⁸ So here are some reasons to think the contrast is in how things sound.⁹

First, the contrast has the familiar effects of changes of appearance. For instance, if you feel a beat under a familiar melody in a familiar way, you will easily recognize the melody. But if you feel a beat in an unfamiliar way, you may fail to recognize the tune.¹⁰ This is what we should expect if feeling a beat not only happens while hearing music, but alters how the music sounds. Second, these influences on perception do not "spread". The beat changes how some experienced processes appear, but not others. And this shows up again in familiar cognitive effects. For example, while hearing a song, listeners can easily and reliably remember whether sounds *within* the song landed before, on, or after the beat they feel. But they will struggle to make such judgments about other perceived events that weren't part of the song—for instance, a dog's distant barking. This suggests that the beat arises not only *while* we hear, but is somehow part of hearing specific processes.¹¹

From here, I assume my contrast is a contrast in how things appear in perception, or what it is like to perceive. My question is: What explains it? On the usual views, there are just three options. The contrast is either a case of conflicting appearances, a change of topic, or a change in mental paint. I'll argue against all three. I'll then motivate and develop my alternative.

These negative points will take time. But they are worthwhile. First, they constrain and motivate my view. Second, my view rejects entrenched assumptions. So I need to explain why we should reject them—especially since I focus on an isolated case study. Finally, my ultimate interest is how my topic bears on theories of perception. To figure that out, we need to look at those theories.

⁸One reason for skepticism is that feeling a beat is not restricted to hearing or to sounds. Also, as we'll see, it is hard to explain what it is like to feel a beat in terms of familiar features of sounds, including their temporal features (such as their rhythm). Thanks to an anonymous referee for pressing this issue.

⁹For additional reasons to think that feeling a beat is perceptual, see Jenny Judge, "'Feeling the Beat': Multimodal Perception and the Experience of Musical Movement", in Peter Cheyne, Andy Hamilton, and Max Paddison, eds., *The Philosophy of Rhythm: Aesthetics, Music, Poetics*, (New York: Oxford University Press, 2019), 76–91, at pp. 79–80.

¹⁰See Peter Essens and Dirk-Jan Povel, "Perception of Temporal Patterns", *Music Perception*, , 4, (Summer 1985): 411–440, at p. 432, and John Sloboda, "The Communication of Musical Metre in Piano Performance", *The Quarterly Journal of Experimental Psychology*, 2, (May 1983): 377–96. Unfortunately, while widely acknowledged, few have tried to measure this effect in a controlled setting. For limited confirmation, see Stefanie Acevedo, David Temperley, and Peter Pfordresher, "Effects of Metrical Encoding on Melody Recognition", *Music Perception: An Interdisciplinary Journal*, 4, (April 2014): 372–86.

¹¹I don't claim we feel a beat *only* when we hear, or that it *only* changes how things sound. As we'll see, neither claim is true. But the beat has good company here: we hear *durations* of sounds, but also experience durations of other events via other senses.

2 Perceptual views

Many philosophers say things appear as they do in any perceptual experience simply because of *what we experience*, or the experience's *subject matter*: the mind-independent things, events, and features (including relations) that we perceive, misperceive, or hallucinate. Let's call this the *perceptual view*.¹²

On the perceptual view, things appear as they do when we feel a beat simply because of what we experience. *Downbeats* and *Upbeats* accordingly differ either by conflicting or changing topic. But I'll now argue that the appearances neither conflict nor change topic. So we should reject perceptual views of the beat.

2.1 No conflict

While you felt a beat under each metronome beep, I felt a beat between each beep. Do the appearances conflict? I'll now argue that they don't.

My task might seem impossible. For we sometimes disagree about where the beat is. If the appearances conflict, the disagreement makes sense. But if not, the disagreement is harder to understand. Later, I'll argue that it is hard but not impossible. Still, for now, I admit that this is a good reason to suspect the appearances conflict.

On the other hand, my argument might seem unneeded. The appearances don't even seem to conflict. At least, my contrast doesn't seem offhand like a change in the sounds.¹³ More, the contrast doesn't seem to fool anyone. No one, for example, checks their speakers. And even suggestible audiophiles, who demand gear for enhancing sound in every conceivable way, don't ask for gear that syncopates everything. Since illusory changes should sometimes fool us, and this contrast doesn't, it isn't an illusory change.

But we can't trust this behavioral evidence. Consider Rubin's face/vase (figure 1). You can see the picture as two black faces over a white background or a white vase over a black background. Seen each way, the picture doesn't seem offhand to change. And no one inspects the picture for

¹²For a representative naive realist version of the view, see John Campbell, *Reference and Consciousness* (Oxford: Oxford University Press, 2002). For a representationalist version, see Michael Tye, *Ten Problems of Consciousness* (Cambridge: The MIT Press, 1995).

¹³Compare to Ned Block, "Attention and Mental Paint", *Philosophical Issues*, 1, (November 2010): 24–63, at pp. 53–54.

moving parts. But our best psychological theories say this change is (among other things) an apparent change in *ordinal depth*, or which parts of the picture appear closer than others.¹⁴ On this view, our sense that the picture doesn't change survives the apparent change: perhaps because of our background knowledge, or a conflicting perceptual state, or something else.¹⁵



Figure 1: Rubin's vase

So, why think *Downbeats* and *Upbeats* don't conflict—why think the beats we feel are compatible? First, I'll narrow down where conflict could arise. I'll then give some reasons to think that it doesn't arise there.

Here is a more detailed preview. Call two series of sounds *soundalikes* if they can't be told apart by perception. And call features *discriminable* if differences in those features can be perceived. There are two possibilities: either soundalikes can differ in beat, or they can't. (Either the beat supervenes on discriminable features of sounds, or it doesn't.) I'll first argue that soundalikes can differ in beat. From here, I'll argue that the beats you and I feel are compatible: a single series of sounds can carry both the beat you feel and the beat I feel.

Why think soundalikes can differ in beat? My reason is that, otherwise, either *Downbeats* or *Upbeats* is a strange illusion. But, worse, so is a lot of musical experience. We should avoid ascribing widespread error. So, absent strong countervailing reasons, we should assume that soundalikes can differ in beat.

¹⁴See Stephen Palmer, *Vision Science* (Cambridge: The MIT Press, 1999). For philosophical discussion, see E. J. Green, "Representationalism and Perceptual Organization", *Philosophical Topics*, , 2, (Fall 2016): 121–48, and René Jagnow, "Ambiguous Figures and the Spatial Contents of Perceptual Experience: A Defense of Representationalism", *Phenomenology and the Cognitive Sciences*, 3, (April 2011): 325–46. For striking examples where the attended figure is seen peering out from *behind* the ground, see Liqiang Huang and Harold Pashler, "Reversing the Attention Effect in Figure-Ground Perception", *Psychological Science*, , 10, (October 2009): 1199-201.

¹⁵For another example where empirical considerations correct our offhand judgments about how things appear, see Ian Phillips, "Afterimages and Sensation", *Philosophy and Phenomenological Research* 87, no. 2 (2013) on afterimages.

To start, return to *Downbeats* and *Upbeats*. Suppose that while you and I felt a beat differently, we heard everything else alike: the same sounds and relative timing, pitch, timbre, loudness, direction, distance, sameness of sound source, and more.¹⁶ Sounds that match in these ways are soundalikes: no one can tell them apart by hearing. So, what I experience is indiscriminable from what you experience. That is, if we make things exactly how they appear to you in *Downbeats* or to me in *Upbeats*, the results are soundalikes: no one, including us, could tell them apart by hearing.

My point is not simply that we heard the same sounds. The point is about how the sounds appear to us. To see the difference, suppose you heard one sound chirp higher than the rest while I didn't. We heard the same sounds. But if we transcribe what we experienced and then play it back faithfully, the results differ discriminably. In this sense, what you experience is discriminable from what I experience. My point is that things are different with *Downbeats* and *Upbeats*. Here if we transcribe what we each experience and play it back, the results will sound exactly alike—neither of us could tell which score was played when. So, if we each represent how things appear to us "by means of an exact copy—and isn't that a good representation of it?—no change shows up".¹⁷

To prevent misunderstanding, I need to pause for an objection. Don't we experience different events: a different series of pulses, the beats we feel? These events can't be actual sounds, since we hear the same ones. But maybe they are hallucinated sounds–a phantom clicktrack.

This thought is miguided. When you start to feel a beat, you simply don't hallucinate sounds no clicks or cowbell seem to chime in. But even if you did hallucinate a pulse, this couldn't settle how you feel a beat. The same problem would rearise: I could hallucinate the same sounds but feel a beat between them, or fail to feel a beat at all. If hearing a metronome doesn't settle how you feel a beat, neither does merely seeming to hear one.¹⁸

So, *Downbeats* and *Upbeats* are experiences of soundalikes that differ in beat. Does that show that soundalikes can differ in beat? No. Either experience or both may be an illusion. But on the

¹⁶Perfect parity is probably impossible; the argument just requires parity up to our powers of discrimination. But I should acknowledge that this kind of parity is probably also rare. Differences in how we feel a beat often cause differences in perceived timing. My claim is just that it is possible for there to be no such differences.

¹⁷Ludwig Wittgenstein, *Philosophical Investigations* (Wiley-Blackwell, 2009), p. 131 (206e).

¹⁸Likewise, it doesn't help to add a second: you can easily feel a beat either under the low or the high alternating beeps. And the same goes for a more phenomenologically apt clicktrack: "faint" rather than "vivid" events, events experienced indeterminately rather than determinately, imagined events, or felt events rather than heard ones. The problem rearises also for Fred Lerdahl and Ray Jackendoff, *A Generative Theory of Tonal Music* (Cambridge: The MIT Press, 1983). They say feeling a beat is associating a "metric structure" with the sounds. But this metric structure is, essentially, a *rhythm* made of timepoints. Unless we add more to the structure, the points above apply.

present view, this illusion is a special kind: it is an illusion of an impossibility. In particular, it is an illusion of a supervening feature, the beat, "budging free" of its supervenience base, the other discriminable features of sounds we hear alike.

Compare to the waterfall illusion. Things alike in position over time move alike: a thing's motion supervenes on its position over time. But if you stare at a waterfall for a minute and then look at a stationary object, it will both appear to move and to stay in one place.¹⁹ This is possible because, although motion supervenes on changes in position, representations of motion don't supervene on representations of position. They sometimes operate independently.²⁰ If soundalikes can't differ in beat, then *Downbeats* or *Upbeats* is like the waterfall illusion: an experience of a supervening feature, the beat, budging free of its supervenience base.

This "result" is dubious. The waterfall illusion is striking and disorienting. The reason is that the way things look is manifestly incoherent. Neither *Downbeats* nor *Upbeats* is disorienting. Why not, if the way things sound is manifestly incoherent in the same way? This view owes us an explanation.

But more importantly, this view counts too much musical experience as illusory. Consider cultural variation in how listeners feel a beat. Certain features of sounds serve as *cues* for the beat. But they serve differently in different cultures. In some West African traditions, for example, sounds that are emphasized in a rhythm typically mark the *final* beat of the measure. But in many European traditions, they mark the start of a measure.²¹ If soundalikes necessarily share a beat, then at least one community is regularly mistaken about where the beat is.²² This is hard to believe. And ascribing widespread error should be a last resort.

So, let's suppose instead that soundalikes can differ in beat. Still, they might have at most one beat. Compare: utterances can match phonetically but differ semantically; still, each utterance has just one semantic value. So, the appearances still conflict in *Downbeats* and *Upbeats*: each series

¹⁹See George Mather, Frans Verstraten, and Stuart Anstis, eds., *The Motion Aftereffect: A Modern Perspective*. (Cambridge: The MIT Press, 1998).

²⁰One reason to suspect this is possible for the beat is that there are also independent cues to the beat-see David Temperley and Christopher Bartlette, "Parallelism as a Factor in Metrical Analysis", *Music Perception*, 2 (2002): 117–49 for an overview.

²¹See Aniruddh D. Patel, *Music, Language, and the Brain* (Oxford: Oxford University Press, 2007), p. 98.

²²We can avoid cross-cultural comparisons and still make the point. The melodic patterns that combine to make most Western music can, like the metronome, be heard in different ways in different contexts by normal Western listeners. For discussion, see my dissertation, *Ways of Perceiving* (Berkeley: University of California, 2021).

of sounds has at most one beat.²³

One plausible view says an artist's intentions or experiences settle where the beat is. Maybe, for instance, the music has the beat the performer or composer feels, or the one she intends the audience to feel, or something else along these lines. Our intuitive judgments suggest as much. In performing "Videotape", for example, Radiohead moves along to the beat that makes the song heavily syncopated. And many think this settles where the beat really is.

But our intentions actually suggest that beats don't exclude, and the appearances don't conflict. For some songs are *deliberately* metrically ambiguous: their authors intend them to carry multiple beats. "Videotape" seems like a good example. Of course, some ambiguity may be compatible with and even parasitic on exclusive distinctions between beats—a kind of musical joke analogous to a garden path sentence, or painting of a duck/rabbit.²⁴ But in some traditions, metric ambiguity is the norm. In many West African traditions, for instance, complex, layered rhythms make multiple multiple meters equally salient.²⁵ And the same goes for Western music inspired by these traditions—for instance, the phase music of Steve Reich. So, metric ambiguity is not always unintended, exceptional, or subversive. Sometimes, it is the norm. So, I conclude that soundalikes can carry multiple beats.

That is my argument that the appearances don't conflict. But hasn't something gone wrong? Again, we argue about the beat. Doesn't that simply show the appearances conflict? No. First, I'll later sketch one way to understand the apparent disagreement. But, second, we seem to be of two minds about the beat. We argue about it. But we also often treat different beats as compatible. At least, music theorists often assume the relation between meter the so-called "musical surface" of discriminable features of sounds is many-one. Likewise, musicians often treat being able to hear a piece of music in a nonstandard meter as a perceptual achievement—like finding a new proof of a theorem, noticing a new symmetry in a geometric figure, or working out a new interpretation of a novel. So, there is room to deny the appearances conflict.

²³Some deny we perceive any features that differ across lookalikes. But on the *rich view*, we can—for example, whether my face is a human face. See Susanna Siegel and Alex Byrne, "Rich or Thin?", in Bence Nanay, ed., *Current Controversies in Philosophy of Perception*, (New York: Routledge, 2017), pp. 59–80. See also Grace Helton, "Recent Issues in High-Level Perception", *Philosophy Compass*, 12 (December 2016): 851–62.

²⁴Compare to "metric fakeouts." For examples, see Justin London's list: https://www.carleton.edu/people/jlondon/downloads-and-media/

²⁵For discussion, see Justin London, *Hearing in Time* (Oxford: Oxford University Press, 2004), p. 85, and David Locke, *Drum Gahu: An Introduction to African Rhythm* (White Cliffs Media Company, 1998).

2.2 No change of topic

If the appearances don't conflict, the perceptual view must say the beats we feel are *different* but *compatible* features of the sounds. Like seeing a thing's color and then feeling its texture, or seeing that the dress is close to you and far from me, *Downbeats* and *Upbeats* differ compatibly by changing topic. But this can't be right.

Some things share your hand's color but lack its texture. These features differ *intensionally*: different sets of individuals across possible worlds (including the actual world) have each feature. This helps make sense of how there is a genuine change of topic when you see your hand's color but not its texture. But no sounds have the beat you feel and lack the one I feel.

To see why, consider a more flexible listener. For any song where you hear the song as syncopated, she can hears the same discriminable features as you—pitch, timbre, loudness, timing, and so on–but hears the song as syncopated or unyncopated. Likewise, when you experience a song in a two-beat meter, she can hear it in a two- or three-beat meter. And so on, for the other ways we can vary how we feel a beat under sounds.

I argued that *Downbeats* and *Upbeats* are compatible. And the argument transposes to your experiences and those of the flexible listener. For any song we pick, that song carries both the beat you feel and the one our flexible listener would feel if she were listening. And, in that case, the beats you and I feel don't differ intensionally—the same actual and possible series of sounds carry both the beat you feel, and the beat I do.²⁶

Some will conclude different beats are not different features at all. But others won't. One natural alternative treats beats as hyperintensionally different *patterns* in the sounds—patterns that can differ even if just the same individuals across possible worlds fit both patterns.

But this view faces two problems. First, flexible listeners are possible. But we aren't so flexible. Once we feel a beat in a new way under a song, we often get stuck that way, and can't easily revert. Why not? If different beats were simply different temporal patterns, and we had already found both patterns in some sounds, this would be mysterious. Attending to one pattern or the other should be easy, having found both. Compare to a 3 X 3 grid (see fig. 2). You can see the dots in columns, or

²⁶Some may say this flexible listener is conceivable but impossible. But we can approximate this flexible listener well enough when hearing the basic patterns that combine to make most songs—drumming patterns, chord progressions, melodies, and so on—that initially seem unambiguous to us. (See my dissertation *Ways of Perceiving, op. cit.* for discussion.) This is good evidence that the flexible listener is possible.

rows. But you don't get stuck seeing just one pattern, even if you can see just one at a time. On the present view, it is mysterious why feeling a beat is any different.



Figure 2: Dots

Second, we can feel a beat in just one way at a time. But we can perceive multiple temporal patterns simultaneously. To illustrate, listen to this *Hemiola* using the link in the note.²⁷ You will hear two series of clicks, high and low. There are three high sounds for every two low ones. Since the durations between sounds in each rhythm are noninteger multiples of the durations in the other, the Hemiola is a *polyrhythm*.

You can hear this polyrhythm in a three-beat meter, with the tactus under high clicks, going by threes. Or, you can hear it in a two-beat meter, with the tactus under low clicks, going by twos. But you *can't* feel a beat in both ways under the polyrhythm at once—you can't hear the song "in three" and "in two" at once.

However, you can experience and attend to *both* rhythms in the polyrhythm at once. For example, with practice, you can tap along to each rhythm with each hand. And if you focus on your hands, you can experience and attend to both rhythms as you tap them out.

Here you experience and attend to multiple temporal patterns, and their complex interrlation, at once. But you can only feel one beat at a time. Why, if feeling a beat is just experiencing or attending to such temporal patterns? The pattern perception view has no good answer.

²⁷"3 over 2 polyrhythm metronome," Drumset Fundamentals, accessed May 24, 2023, https://youtu.be/gthOmTn puls.

3 Entrainment views

Why can we feel a beat differently while experiencing everything else alike? The next family of views says this is because feeling a beat isn't experiencing features of the sounds. Instead, it is coordinating perceived patterns with a patterns in something else: either your *behavior*, *attention*, or *mental paint*. These views differ importantly. But they suffer similar problems, so I'll consider them together.

The relevant views say feeling a beat is a kind of *entrainment*. In particular, a periodic pattern in the sounds, and a periodic pattern in something else, become *phase-locked*. That is, the *phases* of the one pattern—thuds, for example—change in step with the phases of the other—taps, for example. Second, the phase-locking isn't accidental: the period and phase of the second pattern depends causally on the period and phase of the first. But, third, the phases of the second pattern aren't a *reaction* to the sounds—interrupt the sounds, and the associated pattern will keep going, at least for a little while.

I will argue against entrainment views. But to start, I'll focus on views that appeal to patterns of behavior. One version focuses on *actual* behavior—the simple behavioral view. Another focuses on *representations of behavior* or *behavior-guiding psychological processes*—the sophisticated behavioral view.

These behavioral views are well motivated. For there are surprising and interesting connections between feeling a beat and behavior. First, behavioral entrainment is a standard measure of whether and where listeners feel a beat. When an able-bodied listener can't entrain, she often can't feel a beat. Second, getting her to entrain her behavior to music is also a good way to get her to feel a beat if she didn't already—a fact I've exploited several times. Finally, when normal listeners feel a beat, motor representations or processes are often active. Everyone must concede these facts.

However, the simple behavioral view fails for obvious reasons. You can feel a beat and sit still. And not everyone who can feel a beat is able-bodied—paralysis isn't beat blindness. This is probably why no one (as far as I know) endorses this simple view. Even Ryle appreciated these points, so I won't belabor them.

Still, many offer *sophisticated* behavioral views. On these views, feeling a beat isn't entraining *actual* movement to sounds. It is entraining some behavior-guiding psychological pattern to the

sounds. This may be a pattern of motor representation activation, or the deployment of a special kind of know-how, or some other pattern in the representation or processing guiding motion.

I'll now argue that all sophisticated views fail too. The real problem for behavioral views is not that, sometimes, we *don't move* when feeling a beat. The problem is that it wouldn't help if we *did*. For even when listeners *do* behaviorally entrain to sounds, their movements underdetermine how they feel a beat. And if *actually moving* doesn't settle how we feel a beat, representing that movement can't either.

To start, notice that different audiences move differently to the same music. At a classic rock show, most clap on beats one and three. At a gospel concert, most clap on beats two and four. Play the same music for each audience, and they may clap differently while experiencing the same beat and meter. But we don't need two audiences. You can also clap on beats two and four, or beats one and three, without changing how you feel a beat.

Likewise, we can also clap *alike* but feel a beat *differently*. Clap on each beep of the metronome, and you might feel a beat under each clap. Here you're clapping on the tactus. But you can keep clapping on each sound but feel a beat *between* each clap. Now you're clapping on *subdivisions*. Here you move in just the same way to the same sounds, and *entrain* your clapping to those sounds, as before. But you feel a beat differently.

Of course, clapping isn't exceptional. The same goes for bobbing heads, tapping toes, or otherwise getting down. You can change how you feel a beat not only under the sounds you hear, but also under your own synchronized movements. And if getting down in these ways underdetermines how you feel a beat, there is no reason to think *knowing how* to get down or *representing* the same settles it. So, entraining your behavior, or knowing how to, or otherwise representing that behavior, doesn't explain how things sound to you when you feel a beat.²⁸

Before moving on, I want to emphasize the limits of my conclusion. There are undeniably interesting lawlike connections between feeling a beat and behavior. But my question is why things appear as they do when we feel a beat. Behavioral entrainment views say things appear as they do when we feel a beat because of behavioral entrainment. That is the view I've argued against.

On to the remaining entrainment views. On these views, feeling a beat is entraining some

²⁸The problem is not that you can move unconsciously. You can focus on your claps as you move—representing your own motion—but feel a beat in different ways under the represented movement. For simplicity, I omitted this complication, but the argument above could be reformulated to take it into account.

other *mental pattern* to the sounds. But what is the mental pattern? I already ruled out patterns of *hallucinated* events. But another view says the pattern is either in mental events, or in *mental features:* properties of, or relations to, mental events. I'll start with *patterns of attention*, and generalize from there.

The best developed version of this view comes from Justin London²⁹. London says feeling a beat is *attentional* entrainment: entraining a pattern in our *attention* to sounds to a pattern in sounds. More specifically, the pattern is in *peaks* of attention: points in time when we attend *more* to sounds, or pauses, than other times. On London's view, feeling a beat is entraining these peaks of attention to the sounds.

There is a lot more to London's discussion of meter. But his view faces several problems. For peaks of attention need not coincide with the experienced tactus. When you feel a beat under the beeps, you aren't stuck attending more to the beeps than the pauses. You can shift your attention to the pauses, but go on feeling a beat under each sound. Likewise, when you feel a beat under each pause, you can shift your attention to the sounds, and go on experiencing those sounds as landing on subdivisions. Here the distribution of your peaks of attention changes, but how you feel a beat doesn't.³⁰

So, feeling a beat isn't having your attention beat, or entraining its beat to the beating of sounds. But stepping back, this is unsurprising. A succession of experiences isn't an experience of succession, and a changing experience isn't an experience of change. It is no wonder, then, that beating attention isn't attention to a beat.

What about other mental patterns? Does it help to trade beating *attention* for something else? No. It seems to hurt. Consider mental paint. Negatively, these are features of mental states that characterize what it is like to be in that mental state but aren't representational: they aren't the apparent presentation of worldly things, events, or features. But, positively, they are supposed to be *qualities* of mental states modeled on the felt quality of bodily sensations: the distinctive way pains or itches feel, for example. Such bodily sensations may also represent our bodily condition. But the way it feels to be in pain is (allegedly) something else. And this way pain feels is your

²⁹See London, *Hearing in Time, op. cit.*

³⁰The point remains even if we can attend more or less to many things at once. You need not feel a beat under a more-attended event rather than a less-attended one.

sensation's mental paint.³¹

Unfortunately, feeling a beat changes how *sounds* appear. Attending to sounds is directed on the sounds. So we can perhaps understand how attention could change how those sounds appear. But sensations are by design not directed on sounds. So it is unclear why changes in sensation should change how things sound. A throbbing headache, for example, does not change how things sound to you. So why would other periodic bursts of mental paint change how *sounds* appear to you, instead of adding a seemingly unrelated accompaniment, like the headache? Mental paint views leave this entirely mysterious.³²

4 Beats as units

What is it to feel a beat? On my view, feeling a beat is a way of perceiving the timing of events. But the beat is not itself a temporal feature perceived. To start, I'll clarify this general claim by situating it in debates about perception. Then I'll develop a more specific proposal about feeling a beat.

Many philosophers try to explain how things appear in perception, or what it is like for us to perceive, appealing only to what we perceive or the content of perception. And many are convinced that nothing else could possibly explain how things appear.³³ But there is more to psychological theories of perception. Psychologists don't simply list what we perceive or ascribe content. They offer theories of how perceptual systems measure features, the cues used determine what features things have, computational processes that operate over representations to determine further representations, and more. More abstractly, theories of perception explain not only what we perceive, but how we perceive what we do, or our ways of perceiving.

³¹For more on this alleged distinction, see Ned Block, "Mental Paint", in Martin Hahn and Bjorn Ramberg, eds., *Reflections and Replies: Essays on the Philosophy of Tyler Burge*, (Cambridge: The MIT Press, 2003), pp. 165–200, and Tyler Burge, "Origins of Perception", *Disputatio*, 29 (November 2010): 1–38, at p. 15 and p. 29. Not everyone thinks we need mental paint to explain bodily sensations—many think bodily sensation is entirely representational.

³²Some may reply that mental paint serves in modes of presentation of mind-independent features. See, for example, Block, "Mental Paint", *op. cit.*, and Block, "Attention and Mental Paint", *op. cit.*. But how does this work? Mental paint is introduced positively only by analogy with sensations like headaches. So for the reasons I just explained, it is unclear how it could play this role. In any event, this is an open challenge for defenders of mental paint.

³³See Alex Byrne, "Intentionalism Defended", *Philosophical Review*, 2, (April 2001): 199–240, who argues that there can be no difference in how things appear without a difference in content. For a response, see Craig French and Ian Phillips, "Austerity and Illusion", *Philosophers' Imprint*, 20, 15, (May 2020).

These ways of perceiving are no worse understood than what we perceive or perceptual content—in psychology, the two come together. And they are better understood than mental paint—by design, mental paint never appears in psychology (except maybe in more speculative postscripts). So we should be open minded about whether ways of perceiving partly explain how things appear in experience. On my view, they do.

However, these ways of perceiving are a grab-bag. So, unlike other authors, I don't think there is just one theoretically interesting "extra factor" that influences how things appear—for example, the part-whole structure of perceptual representations³⁴, the intentional mode used to represent a content ³⁵ or the point of view of the perceiver. ³⁶ The remedy for the perceptual view is not tacking one more parameter onto our theory. The other factors influencing how things appear are importantly diverse, and may contribute to how things appear in importantly different ways. Recognizing their influence on how things appear is the start of a theoretical project, not the end.³⁷

That is the general approach. And my view is that feeling a beat is a way of perceiving. But more specifically, feeling a beat is a way of perceiving events relative to *periods of time* and *timepoints* that divide the periods. On this view, individual beats are *points of time* that are experienced as the boundaries of periods of time. *Measures* are periods of time whose boundaries are set by beats. And when we feel a beat under sounds, the timing of sounds is experienced *relative to* these timepoints and periods in a distinctive way. I'll soon try to clarify this claim. But to get the rough idea, it will help to start with an analogy.

Suppose you want to represent the timing of the metronome beeps with a pen and paper. How should you do it? One way is to use a timeline or piano-roll notation: you mark horizontal bars, each representing a beep, an even distance apart, each representing a pause. Another way is to use a stopwatch to measure each event and pause in units—say, seconds—and then report your measurements: "one beep, .10 seconds; one pause, .9 seconds;" and so on. But yet another way is to use a *calendar system* to measure the same events and pauses, again in units. So, for instance,

³⁴See Kevin Lande, "The Perspectival Character of Perception", *Journal of Philosophy*, , 4, (April 2018): 187–214, and "Mental Structures", *Noûs*, , 3, (September 2021): 649–77.

³⁵See Adam Pautz, "Experiences Are Representations: An Empirical Argument", in Bence Nanay, ed., *Current Controversies in the Philosophy of Perception, op. cit.*, pp. 23–42.

³⁶See Campbell, *Reference and Consciousness, op. cit.*

³⁷Of course, I'm not the only one attracted to this view. For additional recent arguments that we should take up this project, see French and Phillips, "Austerity and Illusion", *op. cit.*, and Ori Beck, "Rethinking Naive Realism", *Philosophical Studies*, , 3 (2019): 607–33.

you might keep track of each event and pause's timing relative to the *time of day*, and record where each beep and pause falls within the time of day.

The stopwatch and the calendar both represent the timing of events in units, unlike the piano roll. But there are important differences. We use a stopwatch to measure an event from start to finish in units. But even if you can define one day in terms of those same units of time, calendar days do not wait to start until the first event on the calendar day. We use a calendar not just to measure individual events start to finish. We use it to set up a temporal coordinate system of regularly repeating unit durations of time.

On my view, feeling a beat is a way of perceiving the timing of events that is (in some ways) like a special use a calendar system. Rather than using calendar days and so on, the system uses different units of time: beats and measures. And rather than relying on an external medium, we simply perceive events as they happen in terms of the units. Still, we perceive events and their timing relative to a temporal structure characterized by the repeated application of units of time: beats and measures. This is a bit like updating a special kind of calendar "live", as things happen.

The analogy has limits. Just like a calendar's days march on even when the day's events change their pace, the beat you feel sometimes marches on at the same pace even when the sounds change pace. For instance, if you feel a beat under the metronome, and then the beeps speed up, you may go on experiencing their timing relative to the same slower beat as before. Likewise, periods may go unfilled: just as some rare Sunday can lie empty on your calendar, a beat may go on through silence. However, unlike a calendar day's indifference to the day's events, the duration of the beat you feel changes in response to what you hear. The duration of a beat is typically initially scaled in complex ways to the timing of perceived events, and then rescaled as events change. My view leaves open how exactly this happens.³⁸ For instance, if the metronome speeds up for long enough, the beat you feel will typically (eventually) speed up to match. On my view, you're now perceiving the timing of the beeps using different, shorter units. I'll soon offer a hypothesis about why this should happen. For now, I just want to flag the disanalogy.

³⁸The relations between beats and sounds are certainly complicated. Durations between beats don't even typically neatly match the durations between sounds. And the timing of beats doesn't even typically neatly match the onsets of sounds. Instead, beats relate in complex and dynamic ways to several (possibly competing) cues defined over longish series of sounds. (For an overview of the complexities, see Temperley and Bartlette, "Parallelism as a Factor in Metrical Analysis", *op. cit.*) So, if beats are timepoints that separate unit periods of time, those periods of time must relate in those same complex ways to the cues. Fortunately, my view leaves open these empirical issues.

So, like the familiar calendar units, beats and measures need not match durations between events represented in terms of those units. But, unlike familiar calendar units, beats and measures are typically scaled and rescaled in complex ways in response the durations between perceived events. The key point of the analogy, however, is this: when we feel a beat, unit periods of time are applied in an additional structure that is used to perceive the timing of events. And when we feel a beat under events, we not only experience the timing of the events, but also experience their timing relative to these (additional) timepoints and periods. For this reason, the view will be one precisification of the popular idea that feeling a beat isn't simply perceiving the timing of events, but associating a further "temporal pattern" with those events.

That is the basic idea. But now I need to clarify the view. I'll start by clarifying how feeling a beat is a *way of experiencing* the mind-independent subject matter of perception, rather than simply experiencing an exotic subject matter built of units of time. I'll then explain more fully what it is to experience events *relative to* these timepoints and periods. I'll then motivate my specific claims about beats and measures by looking more closely at musical practice. And, finally, I'll develop the view against some important challenges.

4.1 Way of experiencing vs. what is experienced

On my view, we experience beats and measures—timepoints and periods of time. Those units are part of the subject matter of experience, then. So why isn't this an explanation of how things appear in terms of the subject matter of experience—a perceptual view?

One reason is that these periods of time and timepoints are not new mind-independent events, processes, or features. They are *units* that are used to perceive things, events, and features. And while these units are important for characterizing how things sound to you, they don't distinguish ways the mind-independent world seems to be.

But why aren't periods mind-independent? Events happen periodically no matter how we think about, talk about, or experience them. The sun, for example, rises, and sets, the seasons change, the amplitude of a pressure wave oscillates, and a metronome beeps, no matter how we perceive, think, or talk about them. These processes are in this sense objective. And they are also objectively *periodic*: they run through their phases in a (roughly) fixed duration, independently of how we

perceive, think, or talk about them.

Periods of time as they figure in feeling a beat are not objective for a different reason. First, there is no privileged way to represent a process using periods of time. Consider the start and end of each period. Does the sun, for example, *first* rise, and then set; or does it *first* set, and then rise? Do the four seasons really run winter, spring, summer, fall; or spring, summer, fall, winter; or...? And does each oscillation of the wave's amplitude really start at its highpoint, or lowpoint—is it really a sinewave, or a cosinewave? These are bad questions.

Things are different if we *define* periods, which *start* and *end* at specific times, or some phases of a process. In a calendar system, for example, we define several related periods: days, weeks, months, years, and so on. Each starts, and ends, at certain times. So, questions about whether, say, winter comes at the start, or end, of a year make sense. But we could use other calendar systems to represent the same goings on just as well, in different terms.

If we want to represent a periodic process accurately, no calendar is privileged—the "book of the world" need not be written using one calendar rather than another. Periods are a good device for representing the timing of events, but they are not "extra structure" in processes that distinguishes them. And, as a rule, we don't get confused about this. There are few bogus disagreements about which calendar is accurate, or whether a choice of coordinates has mistaken sinewaves for cosinewaves. On my view, beats and measures are like days and weeks: units used in a system of representing time.

This brings us to a second reason my view isn't a perceptual view. Consider the usual explanatory ambitions of the perceptual view (as I defined it). As John Campbell puts it, on perceptual views, the character of perception is "inherited from" the character of the world perceived or represented.³⁹ As I put it, things appear as they do *because of* our experiencing what we do.⁴⁰ And on pain of problematic circularity, the relevant character of the world is independent of how it is perceived, thought about, or otherwise represented by minds or perspectives on the world—the world as it is "anyway", independent of us. But periods of time are not part of the character of the world, so understood. They are not features of the world as it is "anyway", independent of systems

³⁹See John Campbell "A simple view of colour," in John Haldane and Crispin Wright, eds., *Reality: Representation and Projection* (Oxford: Oxford University Press, 1993), pp. 257–268.

⁴⁰For further discussion of inheritance about color experience, see Mark Eli Kalderon, "Color Pluralism", *Philosophical Review*, , 4, (October 2007): 563–601.

for representing that world. They are parts of systems for representing the world. So, feeling a beat does not inherit its character from the character of the mind-independent world experienced. Here the character of a system of representation explains the subject matter of perception, not the other way around.

Does this mean feeling a beat is instead explained by qualia? No. For the relevant system of representation can't be explained in terms of qualia. As we saw, we can't explain how things appear in feeling a beat in terms of mental states modeled on sensations. If we could, that might be a start for explaining the relevant system of representing timing in terms of qualia. But without it, there is no hope of doing that.⁴¹

That is the sense in which beats and measures aren't mind-independent, on my view. Now I need to explain what it is to perceive the timing of events *in terms of* these units.

4.2 Perceiving in terms of units

What is the relevant way of perceiving the timing of events *in terms of* units? This question is pressing because there are ways of perceiving timing that involve units but are nothing like feeling a beat. Maybe, for example, you have a hunch that you've been reading this section for a couple minutes. Here you're aware of a period of time. But you aren't feeling a beat. Nor is feeling a beat simply awareness of special *kinds* of timepoints and periods. For beats or measures often coincide with standard units for timekeeping. You can, for example, feel a beat at 60bpm, or 1 beat per second. If you check your watch, you may become aware that a few seconds have passed, but you aren't feeling a beat.⁴²

If feeling a beat is perceiving events in terms of units of time, why isn't feeling a beat like checking your watch? I think the key difference is that, when you feel a beat, but not when you stare at your watch, you hear the sounds as "organized" by periods of time. So, I'll now try to get rid of those scare quotes—to say what it is to experience events as being organized by periods of time.

⁴¹Here I depart from typical Fregean views of perceptual experience. (See, for example, David Chalmers, "The Representational Character of Experience", in Brian Leiter, ed., *The Future for Philosophy*, (Oxford: Oxford University Press, 2004).) Those views explain how things appear in terms of the subject matter of perception. But the relevant subject matter concerns the normal causes of experiences with certain qualia, or other relations to qualia.

⁴²Likewise, since you can look at your watch to time events as they happen, it doesn't help to add that feeling a beat is experiencing timing "as it happens" using units, rather than retrospectively or prospectively.

My view is that the organization is a special kind of *perceptual grouping*. Specifically, when we feel a beat, we experience events and pauses as going together in larger groups or wholes, *periods of time*. Call this *grouping by period*. To explain this view, I'll need to say a bit about perceptual grouping more generally.

First, return to the dots in fig. 1. You can either see the dots as making rows or as making columns. Focusing on such examples, some say grouping is perceiving *common features*. ⁴³. Here, for example, you either see that the dots *make columns* or that the dots *make rows*. But this view fails.⁴⁴ As E.J. Green points out, if you rotate the page, you see a row become a column. But this doesn't destroy the seen group. Since you see the same items as making the same group, but no longer see them as a row, seeing the group can't be seeing that things make a row.⁴⁵

On Green's view, grouping is perceptually representing things as composing a further thing: a whole. In particular, it is representing those things as forming a whole *and* as sharing certain features—the features that tend to cause us to group items, or *cues* to grouping. In the case of visual grouping, the cues include proximity, common fate (e.g., moving uniformly in the same direction), connectedness (e.g., whether a bar connects the elements), and more. So, in the case of the dots, grouping the dots into columns is representing that they compose a whole and that they share some of those cues.

This view improves on the feature-sharing view. But it has serious problems. For one thing, part-whole relations are cheap. On the usual logic of mereology, composition is unrestricted: any two things make a further whole. But grouping happens only when things present certain cues: proximity, sharing a common fate, and so on. If so, perceptual systems do pointless detective work before they represent a trivial relation between things—a bit like representing a thing as being self-identical only when it is red. Since this would be of little use, we should suspect that the relevant property is not simply the property of making a whole.⁴⁶

A better view treats perceived groups as special kinds of wholes, whose existence the cues

⁴³See, for example, Michael Tye, "Visual Qualia and Visual Content Revisited", in David Chalmers, ed., *Philosophy* of Mind: Classical and Contemporary Readings, (Oxford: Oxford University Press, 2002): 447–56.

⁴⁴For an excellent overview, see Johan Wagemans, James Elder, Michael Kubovy, Stephen Palmer, Mary Peterson, Manish Singh, and Rüdiger von der Heydt, "A Century of Gestalt Psychology in Visual Perception: I. Perceptual Grouping and Figure-Ground Organization", *Psychological Bulletin*, , 6 (November 2012): 1172–1217.

⁴⁵See Green, Representationalism and Perceptual Organization", op. cit..

⁴⁶What about the property of making a whole *and* having some of the cues? As Green notes, the relevant wholes may lack any of the cues and perceiving all of them is not yet perceiving that something makes the relevant whole.

indicate but don't constitute. Membership in the kind takes more than having the cues. And as we will now see, this view can help us make sense of feeling a beat: we can theorize feeling a beat as experiencing special kinds of wholes. My view is that these wholes are periods of time.

To prevent misunderstanding, grouping events into periods of time is different from grouping as discussed in the literature on music. There "grouping" is usually reserved for *phrasal grouping*, or hearing several sounds as going together in *phrases*.⁴⁷ And nearly all authors who write about feeling a beat correctly point out that feeling a beat isn't phrasal grouping.⁴⁸ A quick proof is that we can hear the same phrases over different beats. To illustrate, mutter "boots-and-cats-and", repeating. You can hear the words go together in phrases, each starting and ending on a specific word: [boots-and-cats-and].⁴⁹ These phrases' similarity to a bass-cymbal-snare-cymbal pattern will probably make you feel a beat. But you can feel a beat under those phrases in different ways: either starting on the downbeat or upbeat. For example, if you hear the sounds group into the phrase [boots-and-cats-and], you can feel a beat under each "boots" and "cats", or instead under each "and"-just as you could feel a beat under the beeps or pauses between. Here phrasal grouping is unchanged but feeling a beat changes: the same sounds are grouped into the same phrases, but they relate to the beat in different ways. So feeling a beat can't be grouping sounds into phrases. A second route to the same conclusion is that we need not experience any events at all in order for us to feel a beat—as happens when you go on feeling a beat through a pause in the sounds. By contrast, you can't experience phrases made of nothing but pauses.

Still, feeling a beat is another kind of grouping. It is experiencing wholes, periods of time, which can in turn be built up from other points and periods of time. This view has two major parts: the claim that feeling a beat is a kind of perceptual grouping, and the claim that the relevant groups are periods of time. I'll now motivate each one.

One simple reason to think feeling a beat involves grouping is reflection on how things appear. As I noted earlier, when you feel a beat, you experience sounds as "going together" in measures, or

⁴⁷These phrases are analogous to prosodic phrases in language. Simply due to the stress and timing of spoken words, you may hear several sounds "go together" in phrases, whether or not you understand the language and so parse its semantic or syntactic units. We likewise hear stress-and-timing-based phrases of music. See Patel, *Music, Language, and the Brain, op. cit.*, at p. 108.

⁴⁸See, for example, Lerdahl and Jackendoff, *A Generative Theory of Tonal Music, op. cit.*, London, *Hearing in Time, op. cit.*, p. 90, and Patel, *Music, Language, and the Brain, op. cit.*, p. 106.

⁴⁹You can also easily hear [cats-and-boots-and], with the initial "boots and" preceding the first phrase; if so, just make the necessary changes to the rest of the paragraph.

groupings of beats. Such descriptions are rough. But they are the usual starting point for theorizing about grouping. They were, for example, the starting point for theorizing about visible groups, like the dots: when we reflect on how things appear, we find that we see some of the dots as "going together" in larger wholes.⁵⁰ To be sure, recent work aspires to *measure* each cue's degree of influence on grouping. Still, the usual test of whether grouping happens is a subject's judgment about whether various items go together in a larger whole. The experiments then measure how robustly the cues cause the judgments.

My second claim is that feeling a beat is grouping *by period* of time. This is a more specific claim. Why believe it? The main reasons are abductive: it offers a good explanation of thought, talk, and action involving the beat, and better explanations are unavailable. Of course, this is not conclusive. But, taken together with the problems for competing views, I think it gives significant support for my view.

Let's briefly look at how beats and measures are used in musical performance and downstream in thought and talk about music. We will find that the beat helps us keep time in ways that are strikingly similar to other periods of time.

4.2.1 Units for performance and perception

To kick off a piece, Leonard Bernstein waves his hands, Joey Ramone screams "one, two, three, four!", and Miles Davis snaps. These gestures inaugurate the music, but they aren't part of it. Instead, they count off the beat and meter of the piece to come. The gesture may continue through piece, as with the conductor's waves, or not, as with the snaps. Either way, a beat and meter conventionally associated with the count continues past the count.

The count helps performers coordinate their timing. As the inaugural shot starts the race and a chant sets the pace for the crew's rowing, the count helps the orchestra begin at the same time and carry out the planned sounds and pauses at the right rate. But why feel a beat after the count? The reason is that performers need to stay in sync as the piece wears on. Each member of the orchestra reads just one part of the score. And for each part, the noted events must be timed within a small

⁵⁰Wertheimer did not establish principles of grouping with controlled experiments. He showed figures to his readers which clearly illustrated the phenomenon by inducing it. As others have pointed out, we could view the paper and its reception as an ongoing experiment whose subjects are the readers. Still, the effect is demonstrated by a subject's judgment about how things appear to her.

margin of error—no one can rush or lag more than some small amount over some small amount of time. But the composite of all of these parts must also be executed within a small margin of error in timing. If, for example, you rush within that margin of error, but I lag within it, the result may be unacceptable. The whole orchestra must rush or lag as one. Likewise, even when performers aren't reading from a score at all, they need to start playing in sync and stay there.

How can this be done? Each player will need awareness of how much time has passed since the count. But we aren't very good at estimating large durations between events. Performers need to find another method.

Performers might use a watch or some other standard timekeeper, like the waving of the conductor, to keep time. But perceiving then reacting to an external time-keeper takes time. And, for the purposes of most musical performances, the lag is too great. Performers do better if they keep their own time.

However, musicians also need to keep track of where other performers are. For their goal is usually not for each part to unfold with rigid metronomic accuracy within a margin of error. Deviations from metronomic accuracy will happen. And when they do, they should propagate so that performers rush or lag roughly as one. If all the members of the orchestra are simply keeping track of how much time has passed, and each keeps time imperfectly, some will rush or lag more than others. They will need some way to get feedback about how the other performers are keeping time to correct course.

So, performers need to keep their own time but defer to feedback—to recalibrate. How do they do it? As any performer can tell you, they do it by feeling a beat. And my view can explain why this experience helps. On my view, when they feel a beat, the passage of time is experienced in terms of beats and measures, unit periods of time. This experience plays the role of an external timekeeper without the slow mediating process of perceiving and reacting to such a timekeeper.

The fact that feeling a beat helps us keep time is preliminary motivation for my view. Further support for the view comes from comparing the role beats and measures play to the roles played by other units of measure.

We use units of measure to compare magnitudes—the length of one bridge with that of another, the duration between two events with that between two others, and so on. And to be used as a unit is, among other things, to be used in *quantifying* comparisons between magnitudes: for example, to be used in checking whether one bridge is longer than another—a merely *ordinal* relation between magnitudes—or whether one bridge is twice as long as another—the *ratio* of two magnitudes. Without getting into controversial issues about measurement, we can see that beats and measures function like units in comparisons of temporal magnitudes.⁵¹ Much as ordinary speakers talk about *how many seconds* long something is, we find musicians talking about *how many beats* long a phrase is. And much as we describe the periodic alternation of tides in hours, we find musicians describing cyclical patterns in terms of beats and measures—for example, that the opening chordal motif of "Videotape" is four bars long, that the first chord to sound plays one beat before the start of the song's first full measure, and so on.

There is at least one important difference between beats and other units of time: a second of one process and a second of another are the same duration, but one beat of a song need not be the same duration as one beat of another. However, such variation in units happens elsewhere: we rescale units to suit what we are comparing in order to simplify calculations, to make variance in coordinates reflect only variance of interest in the data, and so on.⁵² Beats and measures are tied to a given piece's periodic patterns much like earth or Jupiter days are scaled to different astronomical patterns.

So, beats and measures seem to play a role in thought and talk that is a lot like the role played by other units. But turning to their role in perception, it is much harder to make a comparative case. For as Christopher Peacocke emphasizes, we don't typically perceive magnitudes in *units of measure*.⁵³ You see my height, for example, but not in inches or centimeters. And even if an expert can reliably estimate those measurements, my height won't look different to her depending on her choice of unit. Even her perception of magnitudes is *unit-free*. The periods of time I'm appealing to would be an exception. The same unit-free distribution of sounds over time is experienced, feeling a beat either way. Then, it is experienced as aligning differently in relation to different periods of

⁵¹See Stanley Smith Stevens, "On the Theory of Scales of Measurement", *Science*, , 2684, (June 1946): 677–680. See also Christopher Peacocke, "Magnitudes: Metaphysics, Explanation, and Perception", in Annalisa Coliva, Volker Munz, and Danièle Moyal-Sharrock, eds., *Mind, Language and Action: Proceedings of the 36th International Wittgenstein Symposium* (Berlin: De Gruyter, 2015), pp. 357–387 for defense of realism about magnitudes.

⁵²One example here is *principle components analysis*, a method for finding coordinate axes that, given some data set, maximize variance in the data.

⁵³See "Perceptual Content", in Joseph Almog, John Perry, and Howard Wettstein, eds., *Themes from Kaplan* (Oxford: Oxford University Press, 1989) Peacocke, "Perceptual Content", pp. 297–330, at p. 300. See also Peacocke's "Temporal Perception, Magnitudes and Phenomenal Externalism", in Ian Phillips, ed., *The Routledge Handbook of Philosophy of Temporal Experience* (London: Routledge, 2017), pp. 213–224.

time-relative to units.

So, it is hard to make a comparative case here. Still, the abductive case above remains strong. And I take it that we shouldn't reject a view of a puzzling phenomenon because it treats the phenomenon as exceptional—that is why the phenomenon is puzzling.

5 Challenges

I'll now close by developing my view in response to some serious challenges.

First, in feeling a beat, some but not all sounds are organized by the beat. For example, if you hear "Videotape" on headphones, and a dog barks nearby, you won't experience the barking as organized by the beat—like other background noise, it will float free of the beat. Why? In particular, if feeling a beat is simply a way of experiencing the timing of events, and you hear both musical events and others as happening at some given time, why doesn't the beat you feel spread to everything you hear at that time?

We might reply that beats and measures are units that are used to perceive only *events in a given process*. For example, just as you may time your laps but not your breaths in seconds using a stopwatch, you time a song but not ambient noise using the beat. There is independent evidence that events are grouped into separate processes relatively early in hearing, so it is not out of the question that temporal units could be applied only within a given process.⁵⁴ Still, the problem rearises. We often see the same events we hear—for example, the crash of the drummer's stick against a cymbal. And often enough, it is manifestly the same event seen and heard. But the beat doesn't organize the seen events. Why not? If the beat organizes events that are part of some given process, or are so perceived, why doesn't the beat organize seen events that we see as being part of that process?

A second attempt says that the beat is not only process-specific but also *modal*: it is a way of perceiving timing that is used only in some specific sense modality, like hearing. Unfortunately, the beat isn't modal. We can feel a beat in both hearing and touch, and in perceiving mixtures of auditory and tactile stimuli that wouldn't separately make us feel a beat. The beat is thus one (and maybe the only) example of intermodal perceptual organization: an organization sustained by

⁵⁴See Albert Bregman, *Auditory Scene Analysis* (Cambridge: The MIT Press, 1990), pp. 9–10. In his terms, grouping by process is "by stream".

integration of felt and heard events when neither feeling nor hearing would be enough.⁵⁵

My (tentative) reply is that feeling a beat is a way of perceiving timing that operates only in hearing, touch, and in the joint use of hearing and touch. It is a *crossmodal* way of perceiving, but not an *amodal* way of perceiving which can be applied whatever modality is in use. There are interesting questions about what the sense modalities could be in order for such limitations to arise. But other arbitrary-seeming limitations within specific modalities are familiar enough to provide a model.

A second problem concerns the triviality of dividing time into periods (much like the triviality of composing wholes from parts). We feel a beat only when specific cues are present. Specifically, psychologists generally agree that, all else equal, (Western) listeners will feel a beat such that beats land on note onsets, land on longer rather than shorter notes, are timed regularly, land at the start of musical phrases, land on harmonic changes, and land at the start of repeating melodic patterns.⁵⁶., On my view, these cues cause us to group events within periods of time. But why? Like part-whole relations, periods of time are plentiful. And any possible partitioning of the events you hear into periods of time would represent the events you hear just as accurately as the partitioning we all converge on. So why do we find perceptual systems doing pointless detective work—checking whether events present certain cues—before choosing such a partition?

My reply is that even if any partitioning of events into periods of time would represent those events just as accurately, not all would represent those events just as well. For representations can be better or worse, given a practical project, for reasons other than accuracy. One virtue of a choice of unit is simply that others adopt it too—that we "standardize" our choice of units, no matter which units we settle on. Another is that we adopt the units that make descriptions (and later coordinations made using the units) simpler or otherwise more useful.

What makes some units of time simpler or more useful than others? Consider virtues and vices of calendar systems. Astronomers in ancient Egypt worked out that a solar year—the time it takes for the earth to revolve around the sun—is about 365 and 1/4 days. But priests refused to add 1/4

⁵⁵See Charles Spence, "Cross-Modal Perceptual Organization", in Johan Wagemans, ed., *The Oxford Handbook of Perceptual Organization* (Oxford: Oxford University Press, 2015) Juan Huang, et al., "Feeling Music: Integration of Auditory and Tactile Inputs in Musical Meter Perception", *PLOS ONE*, 10, (2012): e48496. Spence claims this is the only instance of perceptual organization that "isn't present in any of the constituent sensory modalities when considered individually" (p. 636).

⁵⁶See again Temperley and Bartlette "Parallelism as a Factor in Metrical Analysis", op. cit.

day to the calendar year. David Duncan writes:

this launched the Egyptian calendar on a slow drift across the seasons in a cycle that repeated itself every 1,460 years. Called the Sothic cycle, this flaw was not corrected until the Ptolemaic era in Egypt.⁵⁷

The "flaw" is not that the Egyptian calendar year is ill defined. It isn't. And the problem is not that it fails to match a "joint" in nature—that the Calendar of the World must be kept in solar years. Likewise, the problem is not that Egyptian calendar years are unusable. We could record history just as accurately and make the same plans with the same degree of precision using Egyptian calendar years rather than solar years. (We can describe the Sothic cycle, or the drift of seasons from the Egyptian calendar year, *in Egyptian calendar years*.) The problem is that the unit is *harder to use* than solar years given that the other processes we plan our lives around are in phase with the seasons. Without the quarter day added to a year, the phases of the calendar year quickly get confusingly out of phase with the seasons and other solar-year-pegged phenomena that we plan our lives around. And since we plan around those phases, we do better to choose units that are rescaled to them.

So, at long timescales, there is pressure to adopt some units rather than others. On my view, the same goes for coordinating musical performance and listening at smaller timescales. There are pressures to standardize our measurements, picking the same units as others; there is pressure to rescale units to match the period of the periodic process we are measuring for simplifying calculations and plans; and no doubt there are other practical pressures as well. On my view, the standardization of cues for feeling a beat within a community relieve these pressures. They lead us to feel a beat in the same ways as others, facilitating standardization, and also to rescale beats and measures to match the period of salient patterns in the music, for the sake of simplicity. They do not lead us to feel a beat in these ways because other ways would be *inaccurate*, but because other ways would be (within our community) nonstandard or more complicated.

Finally, these points help demystify the disagreements we started with. At first sight, those disagreements suggest the appearances conflict. But on my view, we can understand them another way. In arguing about the beat, each party attempts to get the other side to adopt a choice of unit that

⁵⁷David Duncan, *Calendar:: Humanity's Epic Struggle to Determine a True and Accurate Year* (Harper Collins, 1999), p. 26

has various benefits other than accuracy. Of course, unlike explicit arguments about, say, which calendar to use, these disagreements don't seem to be *about* units on the surface.⁵⁸) But more covert metalinguistic negotiation about which terms, concepts, or units to adopt arguably happens elsewhere, and it sometimes looks superficially like ordinary disagreements.⁵⁹). Of course, the issue is complicated, and a full discussion beyond the scope of this paper. Still, given that the present view can potentially explain the disagreements, they aren't proof that the appearances conflict.

6 Conclusion

So, feeling a beat one way, rather than another, isn't experiencing our surroundings as being one way, rather than another. But it isn't a subjective accompaniment to perception, either. It is perceiving the timing of events in a certain way: by dividing time into *periods*. And, just as you can use different calendars to record and plan the same patterns of events, you can perceive the same sounds, and their features, as falling within different periods, and use them to plan, attend, or act accordingly. So, our sense that just the same sounds could be heard, as they are, while feeling a beat in different ways reflects an implicit insight into what kind of thing the beat is: the change is in how we are measuring what happens around us, and not an apparent change in the world.

⁵⁸There are indeed such disagreements. See, for example, Elisabeth Achelis, "Calendar Marches On", *Journal of Calendar Reform*, 24 (1954): 91–93.

⁵⁹See David Plunkett, "Which Concepts Should We Use?: Metalinguistic Negotiations and the Methodology of Philosophy", *Inquiry*, , 7-8 (November 2015): 828–74.