

rite, is engaged in so that a certain *heightening* of life, or of consciousness, is attained. In other words, all sound phenomena that are categorizable as music seem aimed at the *enrichment or intensification of experience* via engagement with organized sounds as such. I claim that this is indeed the central core of the music-making intention. It is this that enables us to construe examples from any number of times and places as examples of the artistic-cultural activity we conceive of as music, despite the absence of any limiting intrinsic sonic characteristics beyond that of mere audibility.

An instructive hypothetical example that we would not count as music, and which our definition in its present state just manages to exclude, is the following. Imagine a sequence of sounds devised by a team of psychological researchers, which are such that when subjects are in a semiconscious condition and are exposed to these sounds, the subjects enter psychedelic states of marked pleurability. Such a sequence is not a piece of music; yet it is humanly organized sound for the purpose (arguably) of enriching experience. It does not, however, seek this enrichment through requiring a person's *attention* to the sounds as such. Sound organized for our own good but which does not ask us to listen to or otherwise actively engage with it is not music.

Now that our definition is approaching adequacy we must add one minor qualification, making explicit something only implicit until now: namely, that the organization of sound must be *temporal* organization if the product is to count as musical. What other sort of organization could there be? Well, one can imagine an art in which the point was to produce colorful instantaneous combinations of sounds—i.e., chords of vanishingly brief duration—which were to be savored independently, each in splendid isolation from the next. My intuition is that we would not regard this art as a type of music (though existing musical knowledge and technique would be relevant to its successful practice). It would be the auditory equivalent of jam tasting or rose smelling—the receiving of a sensory impression, sometimes complex, but one for which temporal development was not an issue. Music as we conceive it seems as essentially an art of time as it is an art of sound.

Our complete definition of music would then go roughly like this:

Music = df sounds temporally organized by a person for the purpose of enriching or intensifying experience through active engagement (e.g., listening, dancing, performing) with the sounds regarded primarily, or in significant measure as sounds.

I believe this formulation covers all that it should—e.g., classical music, folk music, party music, avant-garde music, opera, the varied phenomena studied by ethnomusicologists—and nothing it should not (including Muzak).

Some brief observations on the analysis, highlighting its salient features: (1) The analysis accords to music certain intrinsic characteristics, albeit limited ones—to wit, soundingness or audibility, and temporal structure. (2) The analysis is intentionalistic and human-centered; it is people (or the near like) who make music in a purposive way, and not unthinking Nature. (3) The analysis purports to be adequate in application cross-culturally, though of course it is not intended as an analysis of any other culture's concept. (4) The analysis explicitly ascribes a normative attitude to the makers or offerers of music: They necessarily conceive or intend their efforts as worth interacting with. (5) The analysis is creator-oriented or creator-driven: Producers make their production music, through their intentional orientation in bringing it forth, and not the receivers or consumers of such production through anything they might do. . . . I now raise [a further point] about the concept of music. The first is that there is a distinction between what *is* music, and what can be *treated or regarded* as music. One way to ignore this distinction is to claim, in the spirit of John Cage's Zen-inspired reflections, that any and all sounds are music. This is simply false, and the most cogent of Cage's reflections fails to establish it. What Cage shows perhaps is that any sounds

can be listened to *as if* they were music (i.e., attentively, with regard to form, with emotional sensitivity), that one can transform (just about) any sonic environment into an occasion for receptive awareness. It does not follow that all sound events are at present music. The whirr of my blender and the whistling of the wind are not instances of music. They could only be such if they were produced, or proffered, for a certain purpose, as indicated above. But one can *adopt* attitudes toward them which are appropriate for music, with varying degrees of reward.

This Cagean view of music connects to a use of "music" with some currency, in which the word serves as in effect a predicate of experience. The rule of usage is roughly this. If there is a musical experience going on—characterized in some phenomenological fashion or other—then there is music; if not, then the situation is devoid of music. The auditor, through having the right kind of experience, determines whether music is present or is occurring; the source of the sounds experienced in the appropriate way, their *raison d'être*—or even whether they actually exist—is regarded as irrelevant. It should be clear that from my perspective this is a degenerate notion of music, which obscures more than it illuminates, and denies to music several features that I have argued are central to it, namely, sentient origin, artistic intent, and public character. Furthermore, it is a hopelessly relativistic notion, making the status of anything as music (even Mozart piano concertos) relative to each individual listener and occasion. The concepts of a distinctive musical *experience*, or of the hearing of something *as* music, are useful ones, to be sure, but there is little to be gained by collapsing them into the cultural and objective category of music itself.

ONTOLOGY OF MUSIC

■ Ben Caplan and Carl Matheson ■

Types

Let's start with some uncontroversial facts. In 1817–1818, Beethoven composed the Piano Sonata No. 29 in B flat major, Opus 106, which is known as the *Hammerklavier Sonata*. In 1970, Glenn Gould performed the *Hammerklavier* in Toronto. In 1995, András Schiff performed the *Hammerklavier* in New York. We can conclude that there is something—the *Hammerklavier*—that Beethoven composed and that Gould and Schiff performed. But what sort of thing is this?

Perhaps the *Hammerklavier* is identical to one of its performances. But, which one? There are many. It would be arbitrary to pick any one of them, since there is no reason to think that any one of them has a better claim to being the *Hammerklavier* than all of the others do.

Perhaps the *Hammerklavier* is identical to its many performances. This avoids the problem of arbitrariness, but it raises a new problem: how can one thing be identical to many things? The one is one, not many; and the many are many, not one. So they can't be identical.

We need another idea. Philosophers often distinguish *types* and *tokens*. A good way to get a grip on the distinction is to consider some examples. Here's one. How many letters are there in this inscription of the word 'Canada'? On the one hand, there are six letter tokens: the 'C', the first 'a', the 'n', the second 'a', the 'd', and the third 'a'. But, on the other hand, there are four letter types: 'C', 'a', 'n', and 'd'; it's just that one of those letter types—'a'—has three tokens in that inscription of the word. Here's another example. You're in the express checkout lane at the supermarket.

The sign says "10 items or fewer." You have 12 cans of Campbell's tomato soup in your basket. Are you in the wrong lane? On the one hand, if the sign means 10 item tokens, then you are, because you have 12 cans in your basket. But, on the other hand, if the sign means 10 item types, then you aren't, because the 12 cans in your basket are all of the same type.

Perhaps the *Hammerklavier* is a type whose tokens are its performances. This proposal is more promising. It avoids the problem of arbitrariness, since it doesn't identify the *Hammerklavier* with any one of its performances. It also avoids the problem of logical incoherence, since it doesn't identify the *Hammerklavier* with its many performances. And it allows us to explain how the *Hammerklavier* can have multiple performances: the *Hammerklavier* is a type, types can have multiple tokens, and its tokens are its performances, so it can have multiple performances. Let's call this view *the type theory*.

The type theory is pretty widely held. (At one time it was associated with the work of Peter Kivy; and it is still associated with the work of Jerrold Levinson, Stephen Davies, and Julian Dodd.) Those who hold the type theory disagree on a range of further questions, including the following.

- a. Did the *Hammerklavier* come into existence in 1817–1818, when Beethoven composed it?
- b. What role, if any, does the historical context in which the *Hammerklavier* was composed play in distinguishing the *Hammerklavier* from other musical works?
- c. What role, if any, does the instrument that Beethoven specified that the *Hammerklavier* is to be performed on play in distinguishing the *Hammerklavier* from other musical works?

We discuss these questions in the following sections. In the last section, we discuss an alternative to the type theory.

Creation

Some say that the *Hammerklavier* came into existence when Beethoven composed it in 1817–1818. (This is Jerrold Levinson's view, for example.) Those who say that often offer the following argument for their view: in composing the *Hammerklavier*, Beethoven created it; and, in creating it, he brought it into existence; so it came into existence. In reply, most who deny that the *Hammerklavier* came into existence when Beethoven composed it deny that composition is creation. On their view, composition is more like creative discovery. So, for example, when we say "Beethoven composed the *Hammerklavier* in 1817–1818," we don't mean that he literally created anything; rather, all we mean is that he creatively discovered something that already existed. (This is Peter Kivy's and Julian Dodd's view, for example.)

Those who deny that the *Hammerklavier* came into existence when Beethoven composed it often argue that their view fits better with the type theory. After all, according to the type theory, the *Hammerklavier* is a type, and many think that types don't come into existence (either because they exist at all times or because they exist outside of time). If the *Hammerklavier* is a type and if types don't come into existence, then the *Hammerklavier* didn't come into existence either.

Context

As it happens, Beethoven composed the *Hammerklavier* in 1817–1818 and no one else composed a sound-alike musical work—a musical work that sounds exactly like the *Hammerklavier*—a hundred years later. But that's a historical accident. Suppose that Beethoven had composed the *Hammerklavier* in 1817–1818 and that, in addition, someone else who wasn't aware of Beethoven's earlier composition had composed a sound-alike musical work, the *1918 Hammerklavier*, a hundred years later. In that case, how many musical works are there that sound

exactly like the *Hammerklavier*: one or two? The dominant view is *contextualism*, according to which the answer is two, since it's necessary that the *Hammerklavier* is distinguished from other musical works, not just by how it sounds, but also by the historical context in which it was composed, where that context includes at least who it was composed by and when it was composed. Since the *Hammerklavier* and the *1918 Hammerklavier* were composed by different composers at different times, they were composed in different historical contexts. So, according to contextualism, the *Hammerklavier* and the *1918 Hammerklavier* are distinct. (This is Jerrold Levinson's and Stephen Davies's view, for example.) But some reject contextualism in favor of *sonicism*, according to which the answer is one, since it's necessary that the *Hammerklavier* is distinguished from other musical works only by how it sounds. Since the *Hammerklavier* and the *1918 Hammerklavier* are sound-alikes, according to sonicism they must be identical. (This is Peter Kivy's and Julian Dodd's view, for example.)

Contextualists argue that the *Hammerklavier* and the *1918 Hammerklavier* differ in their aesthetic and artistic properties. For example, the *Hammerklavier* is exciting and original in ways that the *1918 Hammerklavier* is not. So, by Leibniz's Law, according to which two things must be distinct if they have different properties, the *Hammerklavier* and the *1918 Hammerklavier* must be distinct.

Sonicists often adopt a "divide and conquer" strategy in reply. They begin by distinguishing *aesthetic* properties like excitingness and eeriness from *artistic* properties like originality and virtuosity. (The difference between the two sorts of properties is that only artistic properties are explicitly about the relation between a musical work and the community or context in which it was composed.) On the one hand, sonicists argue that the *Hammerklavier* and the *1918 Hammerklavier* cannot differ in aesthetic properties like excitingness. For, according to *musical empiricism*, all of a musical work's aesthetic properties can, in some sense, be heard in it. Since the *Hammerklavier* and the *1918 Hammerklavier* are sound-alikes, we can't hear different aesthetic properties in them. So, by musical empiricism, they can't differ in their aesthetic properties. It might seem that the *Hammerklavier* is exciting in ways that the *1918 Hammerklavier* is not, but that appearance must be a mistake, since it's just not possible for the *Hammerklavier* and the *1918 Hammerklavier* to differ in excitingness.

When it comes to aesthetic properties, the issue between contextualists and sonicists thus comes down to musical empiricism. Contextualists start with the intuition that the *Hammerklavier* is exciting in ways that the *1918 Hammerklavier* is not. This leads them to reject musical empiricism. By contrast, sonicists start with intuitions that support musical empiricism. This leads them to reject contextualism. This is a hard issue to settle, since it has to do with the methodological question of which intuitions we should start with. And it's not obvious which intuitions we should start with. As a result, it's not obvious which view we should reject.

On the other hand, sonicists argue that artistic properties like originality are really properties, not of musical works themselves, but rather of composers and their compositional actions. So originality is not a property of either the *Hammerklavier* or the *1918 Hammerklavier* itself; rather, it's a property of Beethoven and his compositional actions or the 1918 composer and her compositional actions. And, even if Beethoven and his compositional actions are more original than the 1918 composer and her compositional actions, it doesn't follow that the *Hammerklavier* and the *1918 Hammerklavier* themselves have different properties, so it doesn't follow by Leibniz's Law that they're distinct.

When it comes to artistic properties, the issue between contextualists and sonicists thus comes down to whether musical works must possess the sorts of properties, like originality, that we ordinarily take ourselves to be attributing to them. Contextualists insist that musical works

must possess those sorts of properties, while sonicists allow that musical works need not have all of the properties that we ordinarily take ourselves to be attributing to them. This is also a hard issue to settle, since it has to do with methodological questions about what role we want musical works to play in our theories and what sorts of access we have to them.

Instrumentation

As it also happens, Beethoven specified that the *Hammerklavier* is to be performed on piano (on “hammer-keyboard” or “*Hammerklavier*”), and no one else composed a sound-alike musical work and specified that it is to be performed on Perfect Timbral Synthesizer (PTS), an electronic keyboard that can duplicate the timbre of any actual instrument. But that, too, is a historical accident. Suppose that Beethoven had composed the *Hammerklavier* and specified that it is to be performed on piano and that, in addition, someone else who wasn’t aware of Beethoven’s composition had composed a sound-alike musical work, the *PTS Klavier*, and specified that it is to be performed on PTS. In that case, how many musical works are there that sound exactly like the *Hammerklavier*: one or two? According to *instrumentalism*, the answer is two, since it’s necessary that the *Hammerklavier* is distinguished from other musical works, not just by how it sounds, but also by the instrument that its composer specified that it is to be performed on. Since the composers of the *Hammerklavier* and the *PTS Klavier* specified that they are to be performed on different instruments, according to *instrumentalism* the *Hammerklavier* and the *PTS Klavier* are distinct. (This is Jerrold Levinson’s and Stephen Davies’s view, for example.) But sonicists would say that the answer is one, since the *Hammerklavier* and the *PTS Klavier* are sound-alikes.

Instrumentalists argue that the *Hammerklavier* and the *PTS Klavier* differ in their aesthetic and artistic properties. For example, the *Hammerklavier* is thundering and original in ways that the *PTS Klavier* is not. So, by Leibniz’s Law, they must be distinct.

Sonicists can adopt a “divide and conquer” strategy in reply here, too: on the one hand, the *Hammerklavier* and the *PTS Klavier* can’t differ in aesthetic properties like thunderingness, since they’re sound-alikes and musical empiricism is true; and, on the other hand, although there might be a difference in originality, that difference is a difference between composers and their compositional activities, not between the *Hammerklavier* and the *PTS Klavier* themselves.

Contextualism and instrumentalism are independent challenges to sonicism. Suppose that you start with the view that the *Hammerklavier* is a type that is individuated, as the sonicist says, entirely by how it sounds. If you’re then persuaded by the arguments in favor of contextualism, you might come to modify your view and hold that the *Hammerklavier* is a type that is individuated, not just by how it sounds, but also in part by the historical context in which it was composed. If you’re then persuaded by the arguments in favor of instrumentalism, too, you might come to further modify your view and hold that the *Hammerklavier* is a type that is individuated, not just by how it sounds and the historical context in which it was composed, but also by the instrument that Beethoven specified that it is to be performed on.

Wholes

In one form or another, the type theory is widely accepted. But there is a problem with it. (This objection comes from Guy Rohrbaugh.) The *Hammerklavier* is *modally flexible*. Modality goes beyond how things actually are and encompasses how they could have been or how they must be. In this case, the modal flexibility of the *Hammerklavier* is that it could have been different than it actually is. For example, in composing the *Hammerklavier*, Beethoven could have called for a different note here or there. Had he done so, he wouldn’t have composed a different work; rather, the same work—the *Hammerklavier*—would have been slightly different. In that case, something

that is a performance of the *Hammerklavier* in the actual world (where Beethoven calls for certain notes) might not be a performance of the *Hammerklavier* in another possible world (where Beethoven calls for certain slightly different notes), even if nothing about the performance itself has changed. But types aren’t modally flexible in this way. That is, something that is a token of a type in the actual world must be a token of that type in another possible world if nothing about the token itself has changed. Since the *Hammerklavier* is modally flexible but no type is, it follows that, contrary to the type theory, the *Hammerklavier* is not a type.

Perhaps, instead of being a type that has performances as tokens, the *Hammerklavier* is a whole that has performances as parts. (More precisely, perhaps the *Hammerklavier* is a whole such that each of its performances is a part of it and every part of it has a part in common with one of those performances.) On this view, the *Hammerklavier* is a temporally extended event that includes other events—performances—as parts. Let’s call this *the whole theory*. Like the type theory, the whole theory avoids the problems of arbitrariness and logical incoherence, since it doesn’t identify the *Hammerklavier* with any one of its performances or with its many performances either. (Although the whole view does identify the *Hammerklavier* with a single whole composed of many performances, it doesn’t identify the *Hammerklavier* directly with the many performances themselves.) Like the type theory, the whole theory allows us to explain how the *Hammerklavier* can have multiple performances: the *Hammerklavier* is a whole, wholes can have multiple parts, and its performances are among its parts, so it can have multiple performances. And, unlike the type theory, the whole theory can allow for the *Hammerklavier*’s modal flexibility, since wholes are modally flexible. Although it might not be obvious, something that is a part of a whole in the actual world might not be a part of that whole in another possible world, even if nothing about the part itself has changed. For example, something might be a part of your car in the actual world but not in another possible world (in which the part has been removed, say), even if nothing about the part itself (other than its location) has changed. Similarly, something can be a part of one whole—the *Hammerklavier*—in the actual world (where Beethoven calls for certain notes) but not in another possible world (where Beethoven calls for certain slightly different notes), even if nothing about the performance itself has changed.

So the whole theory might allow for the *Hammerklavier*’s modal flexibility. But, the type theorist might say, there is still a further problem with the whole theory. (This objection comes from Julian Dodd.) The *Hammerklavier* is *hearable*. In particular, you can hear the *Hammerklavier* itself—the musical work—by hearing a performance of it. But, on the whole view, when you hear a performance of the *Hammerklavier*, the performance that you’re hearing isn’t the *Hammerklavier* itself; rather, it’s just a part of that whole. Since you can hear the *Hammerklavier* in ways that you can’t hear a whole that has performances as parts, the *Hammerklavier* is not a whole that has performances as parts. So the whole theory is false.

In reply, the whole theorist can point out that she is not alone in having to explain the *Hammerklavier*’s hearability. The type theorist has to explain that, too. And, on the type theory, when you hear a performance of the *Hammerklavier*, the performance that you’re hearing isn’t the *Hammerklavier* itself; rather, it’s just a token of that type. So it seems that the *Hammerklavier*’s hearability is just as much of a problem for the type theory as it is for the whole theory. And, more importantly, perhaps the whole theorist can explain how you can hear the *Hammerklavier* itself by hearing a performance of it. The whole theorist might say that, when you go to a performance of the *Hammerklavier*, you get to hear both a performance of the *Hammerklavier* and the *Hammerklavier* itself. In fact, you get to hear the *Hammerklavier* itself precisely *because* you get to hear a performance of it. We often think that someone perceives both a part of something and that thing itself. In fact, we often think that someone perceives something precisely *because* they perceive a part of it. For

example, when a friend is approaching you, you see the front surface part of your friend. But you also see your friend herself. And it's precisely *because* you see the front surface part of your friend that you see your friend herself. So maybe the whole theory can allow for the hearability of the *Hammerklavier* after all. If so, then the whole theory would remain a promising alternative to the type theory.

MAKING TRACKS: THE ONTOLOGY OF ROCK MUSIC

■ Andrew Kania ■

Philosophers of music have traditionally been concerned with the problems that Western classical music raises, but recently there has been growing interest both in non-Western music and in Western musical traditions other than classical. Motivated by questions of the relative merits of classical and rock music, philosophers have addressed the ontology of rock music, asking if the reason it is held in lower esteem by some is that its artworks have been misunderstood to be of the same kind as classical musical works.

In classical music, the production of the sound event to which the audience listens is the result of two quite distinct groups of actions. First, the composer creates the work by writing a score. Then, a performing artist or group of artists performs the work, of necessity producing an interpretation of it. Often, the composer is closely involved in at least the first performance of a new work, but even then his or her contributions as a composer are clearly distinguishable from those made as a performer.

Two Competing Ontologies of Rock

In *Rhythm & Noise*, Theodore Gracyk argues that rock music is the tradition that has cut out the performing middleman and delivers music straight from the composer to the audience. Gracyk is talking about rock music in a broad sense: not as a style (rock as opposed to heavy metal), but as a wider artistic tradition (rock as opposed to classical or jazz). The thesis that Gracyk develops through the first half of his book is that the primary work of art in rock music is not a "thin" sound structure to be instanced in different performances, as in classical music, but the almost maximally "thick" sound structure encoded on a recording and properly instanced through playback of a copy of the recording—what I will call a "track." He argues for this view in part by providing a history of the rock tradition beginning with Elvis Presley's early recordings at Sun Studios, and hitting its stride with Bob Dylan's first electric albums and The Beatles's shift of focus from live shows to the recording studio.

In *Musical Works and Performances*, Stephen Davies criticizes Gracyk's view, pointing to important rock practices that Gracyk ignores or sidelines, particularly the importance placed on live performance skill in the rock world. In summary, Davies says:

The facts are these: more groups play rock music than ever are recorded; almost every recorded group began as a garage band that relied on live gigs; almost every famous recording artist is also an accomplished stage performer; [and] although record producers are quite rightly acknowledged for the importance of their contribution, they are not usually identified as members of the band.

From "Making Tracks: The Ontology of Music," by Andrew Kania *Journal of Aesthetics and Art Criticism*, Vol. 64, Issue 4, Autumn 2006, pp. 401–414. Copyright © 2006 by The American Society for Aesthetics. Reprinted by permission of John Wiley & Sons and the author.

Elsewhere, Davies also points to the fact that cover versions and remixes are treated more like new interpretations of existing works—more like performances—than like new works in their own right. Davies proposes an alternative account of rock ontology intended to correct these shortcomings. He argues that rock works, like classical ones, are created for performance, but whereas classical works are *works-for-live-performance*, rock works are *works-for-studio-performance*, where works for studio performance implicitly include a part for producer and sound engineers.

I am sympathetic with Davies's reclamation of the importance of live performance skill for rock; however, I believe we can find a place for such values in rock without recourse to the notion of a work-for-studio-performance. Several of the problems with Davies's account of rock come from a tension between the idea of rock works being for-studio-performance and the very rock practices he highlights in his criticisms of Gracyk.

First, although many garage and pub bands may hope to be recorded one day, it is not clear that they write their songs with a part for a sound engineer even implicitly in mind. When playing in the garage or pub, without those technicians, these bands seem to think they are providing audiences with fully authentic performances of their songs, not with performances missing a part. Of course, even pub bands use amplification, so one might argue that the role of engineer is being played by someone, even if that someone is the bass player who also does the sound check at the beginning of the gig. But this much engineering is merely the result of using electric instruments. Live performances of classical works involving electric instruments, from Anthony Ritchie's concerto for amplified acoustic guitar (referred to by Davies), through the weird innovations of the early twentieth century such as the theremin and ondes martenot, to the wind machine in Vaughan Williams's *Sinfonia Antartica*, require an engineer one way or another. That does not make those engineers performers of the works (despite there being an "implicit part" for an engineer to "play"); nor does it make those works ontologically for-studio-performance.

Second, Davies maintains both that rock songs are works for studio performance and that "works for studio performance . . . cannot usually be played live." Any account of rock music that makes live concerts an unusual phenomenon is surely misguided. At rock concerts, even by bands that have produced studio albums, both the musicians and the audiences suppose simply that those bands really are performing their songs live.

This intuition is admittedly defeasible in the face of a theory with more explanatory power, but Davies thinks that his account fits our intuitions about live rock better than mine (to be outlined below). This cannot be so if it virtually rules out live rock shows.

Davies has suggested that rock musicians and fans might be acquiescing in the inferior simulations of recordings that go on at rock concerts simply as the result of current technological shortcomings. More and more equipment is making the move from the recording studio to the stage as its size decreases and its flexibility increases. Perhaps one day all that is achievable in the studio will be achievable onstage. At that point there will be no reason to withhold the label "studio performance" from "live" rock concerts.

There are three relevant responses to this suggestion:

1. As noted above, although rock musicians may use on stage some of the same technology they use in the studio to produce the same sounds, they are still expected to *perform* their songs. There is already technology available to reproduce the sound of a recording on stage—a CD player will do that—but rock audiences want to hear musicians play their instruments and sing, just as do classical and jazz audiences, as the occasional lip-syncing scandal shows.