Motivic Perception and Modularity
Author(s): David Temperley
Reviewed work(s):
Source: *Music Perception: An Interdisciplinary Journal*, Vol. 13, No. 2 (Winter, 1995), pp. 141-169
Published by: University of California Press
Stable URL: http://www.jstor.org/stable/40285692
Accessed: 15/05/2012 07:05

Music Perception
Winter 1995, Vol. 13, No. 2, 141–169

Motivic Perception and Modularity

DAVID TEMPERLEY
Columbia University

There is an important distinction to be drawn in the way different kinds of motivic relationships are perceived. Some relationships are detected quickly and automatically; other kinds are detected (if at all) only slowly and deliberately. There is a phenomenological difference here as well. These differences are nicely accounted for by Jerry Fodor's theory of modularity. It is argued that certain relationships are perceived in a "modular" fashion, and others are not. It is hypothesized that the relationships perceived in a modular way are those between segments that are (a) related by tonal transposition and (b) parallel relative to the metrical structure. This view accounts for the differences between perception of different kinds of relationships and also sheds light on metrical structure in general, the "rehearing" problem, and the issue of "mandatoriness" in musical perception.

Parallelism

In teaching untrained undergraduates about music, as many music theorists have occasion to do, one naturally looks for aspects of musical structure that can be easily perceived. We discuss orchestration and texture, with reasonable success; we discuss duple and triple meter, with somewhat more difficulty; we discuss keys and modulations with caution, knowing that a large portion of the class will probably not perceive them or at least will be unable to articulate their perceptions. One of the most widely and easily heard aspects of musical structure is surely parallelism: the similarity of intervallic pattern between short, closely juxtaposed melodic segments. Probably every student is able to hear the similarity between the first four notes and the second four notes of Beethoven's 5th Symphony, for example (Figure 1), or between the first and second four-measure phrases of Mozart's 40th (Figure 2). Such similarities would seem to be among the most basic and incontrovertible facts of musical perception.

Requests for reprints may be sent to David Temperley, Department of Music, Columbia University, 116th & Broadway, New York, NY 10027. (e-mail: dt3@cunixa.cc.columbia.edu)

Fig. 1. Beethoven, Symphony no. 5, I, mm. 1–5.

Fig. 2. Mozart, Symphony no. 40, I, mm. 1–9.

The psychological reality of parallelism has been experimentally demonstrated by Diana Deutsch (1982, pp. 304–311). Deutsch played subjects 12-note sequences and asked them to write them down.[1] Some were highly structured sequences such as Figure 3a, with repeating parallel fragments. Others were unstructured, such as Figure 3b. Not surprisingly, people recalled the structured sequences with much more accuracy. This points up an important fact about parallelism: It is an aid to memory, allowing us to encode what we hear in a parsimonious and efficient way. The ubiquitous parallelisms and sequential patterns in Christmas carols, popular songs, and the like therefore serve an important function: They allow these melodies to be easily learned. In another experiment, Deutsch played listeners only the structured sequences, but this time with pauses inserted, as shown in Figures 4a and 4b. When the pauses occurred between parallel groups, as in Figure 4a, performance on the recall test was not affected. But when

a.
b.

Fig. 3. From Deutsch (1982).

a.
b.

Fig. 4. From Deutsch (1982).

1. The subjects here were trained musicians; unfortunately I know of no comparable experiment using untrained listeners.

they overlapped groups, as in Figure 4b, performance deteriorated dramatically. This experiment shows the importance of rhythm in parallelism. In identifying and encoding parallel fragments, we do not consider pitch alone; if the rhythmic structure reinforces the parallelism, it will be much more easily recognized.

Parallelism is not an easy word to define. In referring to patterns of intervallic similarity, it is clearly related to motivic or thematic structure. But not all motivic relationships are normally considered parallelisms; rather, the term usually refers to relationships that are particularly obvious or audible. In part, my aim in this paper is to examine what it is that makes certain motivic relationships so much more obvious and audible than others. In addressing problems of motivic structure, one is of course faced with a huge and diverse body of work in music theory, ranging from the Grundgestalt approach of Schoenberg and his followers (Rudolf Reti, David Epstein, and Walter Frisch), to the Schenkerian conception of motive favored by theorists such as Carl Schachter and John Rothgeb, to the more ad hoc approaches of Donald Francis Tovey and Charles Rosen. Semiotic analysts such as Nicolas Ruwet and Jean-Jacques Nattiez have explored motivic structure as well, treating it as an aspect of "paradigmatic relationships." Leonard Meyer considers motivic structure in terms of what he calls "conformant relationships"; a conformant relationship, for Meyer, is where "one...discrete musical event is related to another such event by similarity" (1973, p. 44). Finally, motivic structure is discussed in the generative theory of Fred Lerdahl and Ray Jackendoff, who use the term parallelism but also refer to "associational structure," a network of motivic segments related by similarity.[2] These authors differ widely in their approaches to the issue, most basically in terms of purpose. Lerdahl and Jackendoff are interested in motivic structure as it affects listeners' spontaneous hearings of pieces. The purpose of Schoenberg, Reti, and others, on the other hand, is surely more prescriptive; they are pointing out relationships that are not spontaneously perceived, but that might enrich the musical experience once the listener is aware of them. In this paper, I will be approaching the issue from an entirely descriptive point of view: What kinds of structures do listeners spontaneously hear, without a score and without analysis? How might the motivic structure of a piece, as heard by a listener, be described?

2. For examples of work in the Schoenbergian tradition, see Schoenberg (1950) (especially the essay "Brahms the Progressive"), Reti (1951), Epstein (1980), and Frisch (1984); for examples of Schenkerian work on motive, see Rothgeb (1983) and Schachter (1983); for a contrasting approach, see Tovey (1939) and Rosen (1971). For a discussion of relevant work in semiotics (most of which is available only in French), see Monelle (1992), especially chapter 3 (on Ruwet) and chapter 4 (on Nattiez). Meyer's discussion of conformant relationships is found in *Explaining Music* (1973), chapter 3. For Lerdahl and Jackendoff's discussion of associational structure, see *A Generative Theory of Tonal Music* (1983), pp. 16–17 and 286–287.

In this sense, my purpose dovetails most closely with that of Lerdahl and Jackendoff. Ultimately I hope to tie my conclusions in with their theory and to offer some ideas about the nature of "associational structure," which they themselves leave largely unexplored. It should be kept in mind that I am particularly concerned here with low-level or local motivic relationships: those that occur within or between phrases, rather than between large sections of a piece (I will elaborate on this distinction later).

The distinction between the descriptive and the prescriptive approaches to analysis—between describing the way people hear spontaneously and exploring new ways of hearing—is an important one here, because it is clear that the motivic relationships discussed by analysts vary enormously in the ease with which they may be heard. It is surely uncontroversial, for example, that inversions and retrogrades, often discussed in analyses of serial music, are more difficult to hear than simple transpositions such as the one cited in Mozart's 40th. There are important questions here about the nature and extent of this difference, and the basis for it—whether it is due simply to cultural conditioning or to innate capacities of the ear and mind.[3] I will return to these questions later. We should also note, however, that even relationships of simple transposition vary greatly as to their audibility. The importance of rhythm in this regard has already been mentioned; perhaps even more important is the role of meter. When two phrases are juxtaposed that are alike in rhythm and intervallic pattern and similarly placed with respect to the metrical structure, as in the Mozart, their similarity seems to leap out at us in an utterly automatic, immediate way. But if melodic segments differ with respect to the metrical structure, they become much less readily perceptible. Consider Figure 5, the beginning of the first movement of Haydn's String Quartet op. 76 no. 1. This passage contains, in a sense, an extremely strong parallelism: the second three notes of the cello melody are a simple tonal transposition of the first three (one could argue that two other variants of the motive also occur in the phrase; but let us consider only the first two occurrences). However, this parallelism is

3. A number of experimental studies have explored listeners' ability to identify melodic transformations of various kinds: repetitions, transpositions (exact and tonal), and contour-preserving variants, as well as inversions and retrogrades. To my knowledge, no study has been done that specifically compares the perceptibility of transpositions with that of inversions and retrogrades. For a summary of experimental work on the perception of melodic patterns in general, see Dowling and Harwood (1986), pp. 130–144. For a discussion of experimental work on the perception of serial transformations in particular, see Krumhansl, Sandell, and Sergeant (1987), pp. 51–52.

The perceptibility of different kinds of pitch relationships has also been discussed, more speculatively, by music theorists. See Meyer (1967), pp. 266–293, and Browne (1974), pp. 395–401, for arguments against the perceptibility of serial relationships. Schoenberg (1950, pp. 107–114) seems to argue for the perceptibility of serial relationships on the basis of an analogy with vision; more recent arguments in favor are difficult to find, but see Morris (1987), pp. 233–237, 299.

Fig. 5. Haydn, String Quartet op. 76 no. 1, I.

only weakly perceived, if at all; I feel that I was in no sense aware of it the first few times I heard the piece. The reason for this, I submit, is that the parallelism goes against the duple meter of the passage (established by the first three chords, and also reinforced by the two-note grouping of the articulations). Even when one knows it is there, it hardly seems to leap out at one in the same way as, for example, the relationship between the two dyads B-G and C-A in measure 3. The importance of meter in this example becomes clear if we present the melody, with exactly the same pitches and rhythms, as being in $\frac{3}{4}$ time instead. (This we can do by simply changing the opening chords, and by changing the articulations to suggest a three-note grouping, as shown in Figure 6; the melody is now heard, I suggest, with a strong $\frac{3}{4}$ meter). Now, reinforced by the meter, the parallelism leaps out at us very strongly; it becomes a direct fact of musical perception that, I submit, even untrained listeners would notice without difficulty.

In other cases, a pattern of parallelism may be insistent enough to become perceptible even when it is going against the metrical structure. In most of these cases, however, it seems to me that what is really happening is that we are creating—perhaps only momentarily—a competing metrical structure to accommodate the parallelism. The passage shown in Figure 7, from the fourth movement of Beethoven's String Quartet op. 59 no. 3, is

Fig. 6.

Fig. 7. Beethoven, String Quartet op. 59 no. 3, IV.

illustrative. We do hear the three-eighth-note pattern in the first violin line (in mm. 60–62); but in hearing this pattern, we create a secondary metrical pattern to go with it, in which every third eighth-note beat is strong. This recalls Lerdahl and Jackendoff's observation that parallelism is a factor in the determination of meter; we will prefer a metrical structure in which motivic parallelisms are reinforced (1983, pp. 74–75).

It is surely no secret that the perception of motivic relationships is affected by meter, and that segments that are parallel with respect to the meter—we might call them "metrically parallel"—are more readily heard as being related. It seems clear, for example, that there is an experiential difference between the two versions of the Haydn (Figures 5 and 6), in terms of the parallelism in the cello melody. It is not so easy, however, to put this difference into words, in conventional terms of folk psychology. It is clearly not a matter of our knowledge or beliefs about the relationship. You knew the relationship was there even before you heard the passage shown in Figure 6. One might say that the difference relates to what we are able to detect through listening alone. If (assuming you did not know the piece) I had simply played you the passage shown in Figure 6 and asked you if there was a close relationship between the first three notes of the melody and the second three, you could undoubtedly have told me quite easily that there was. If I had played you the passage shown in Figure 5 and asked you this, it would have been more difficult. But it seems to me that this is not really the point. A well-trained musician, could, with little effort, detect the relationship in Figure 5 simply by listening, if only by identifying the intervals one at a time and comparing them. This brings us, of course, to the question of what kind of relationships can or cannot be heard and what it means to "hear" a relationship. If hearing a relationship means being able to detect its presence through listening alone, then it seems to me that *anything* can be heard, even very complex relationships such as those used in serial music. Consider, for example, the clarinet melody from the theme of Webern's Symphony op. 21, shown in Figure 8 (admittedly a relatively simple example). It would not be easy to verify through listening that the second nine notes were a transposed retrograde of the first; but with enough time and effort, it certainly could be done. By this definition of "hearing," I would argue, we can essentially hear any relationship that we can conceive of.[4] In seeking to describe the experiential difference between the two versions of the Haydn melody, then, the distinction between relationships that are detectable through listening and those that are not is not a particularly useful one. Alternatively, we could simply say that there are different ways of acquiring knowledge: One may come to know something in a perceptual way or in a more indirect, inferential way. It is clear that a distinction of this kind can be made in many situations. Seeing someone walk in with a wet umbrella is different from actually looking out the window and seeing it is raining. It seems odd to make such a distinction here, however. In both Figure 5 and Figure 6, we are detecting the relation-

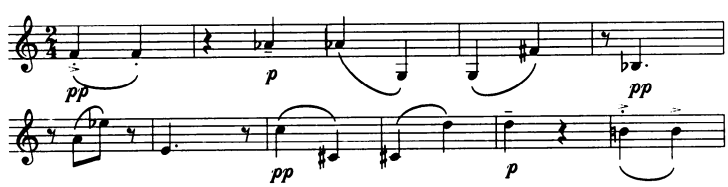

Fig. 8. Webern, Symphony op. 21.

4. This is not strictly true. We cannot for example, hear patterns among pitches that are too high in frequency for us to hear at all. More precisely, then, we could say this: if we are able to identify a set of musical units, we can detect any pattern that we can conceive of that is defined in terms of those units.

ship (if we do detect it) through what would seem is an entirely direct manner: by listening to the music. On the face of it, there is nothing indirect or inferential about detecting the relationship in Figure 5. Why is it, then, that the two examples are experientially so different?

In short, the experiential difference between Figure 5 and Figure 6, with regard to the first six notes of the melody, is difficult to describe. It is not a matter of our beliefs: we know the relationship is there in both cases. It is not a matter of what we can detect through listening: we can detect it through listening in either case. Describing one case as direct, and the other as inferential, seems unsatisfactory. I wish to propose an alternative solution to this problem. In so doing, I believe it is useful to draw on an idea from recent philosophy and psychology, the idea of modularity.

Modularity

In *The Modularity of Mind* (1983), Jerry Fodor offers a highly influential view of the nature of perception. Perceptions are our direct experiences that things look, sound, or feel a certain way; beyond this, what perceptions are is exactly what is at issue. Fodor believes that perception represents not a final picture of our knowledge about a situation, but only a first stage. Perceptual representations are simply hypotheses about the outside world, which are formed on the basis of limited information and which may be sharply at odds with what we actually believe. This is in contrast to the "plasticity" view of perception, which holds that perception may be influenced without limit by our higher knowledge.[5] Although there is perhaps room for many positions between these two extremes, my own feeling is that Fodor's view contains a great deal of truth and is particularly enlightening when applied to musical perception. In this paper, I will suggest that the modularity theory not only provides a useful theoretical framework for discussing motivic perception in music, but also finds in it a striking and valuable confirmation. First, it is necessary to examine Fodor's view in a little more detail.

Fodor's aim in *The Modularity of Mind* is to explain what it is that distinguishes perceptual processes—which for him include not only vision, hearing, and the other senses, but low-level language perception as well—from so-called "central processes," such as reasoning and problem-solv-

5. For a presentation of the "plasticity" view, see Paul Churchland, *Scientific Realism and the Plasticity of Mind* (1979). Since then, the two authors have debated the issues further: see Churchland's article "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor," in the collection *A Neurocomputational Perspective* (1989), and Fodor's articles "Observation Reconsidered" and "A Reply to Churchland" in the collection *A Theory of Content and Other Essays* (1990).

ing. According to Fodor, perceptual processes are performed by a set of special-purpose systems called "modules"; the role of these modules is to form perceptual representations, which are then evaluated by the central system. Fodor's theory can well be viewed as a set of claims about how these representations are formed and what they are like. Perceptual representations, Fodor argues, are formed in a fast, mandatory manner; we cannot choose to see a visual scene as anything but a set of objects, and we cannot choose to hear an English sentence as merely a series of nonlinguistic sounds. In terms of speed, it takes us well under a second to analyze a visual scene or identify an English word. Furthermore, modules are informationally encapsulated. This means that the information on which perceptual representations are based is extremely limited and is only a small part of the knowledge we actually possess. The evidence for this is seen in optical illusions such as the Muller-Lyer illusion shown in Figure 9. Here the two lines appear to be different lengths, but in fact they are the same length. For Fodor, the interesting fact is that even when we know they are the same length, they continue to look as if they are different lengths. This, then, is an argument for the existence of perceptual representations that are both distinct from our beliefs and to a large extent unaffected by them. Fodor's argument, then, is that these characteristics—speed, mandatoriness, encapsulation, and several others that I will not discuss here—apply to perceptual processes but not to central processes. Examples of a central process would be solving a detective story, playing chess, or doing a complex math problem. These are slow processes, and they are optional; we can decide whether or not to do them. More importantly, they are unencapsulated. This means that, in reaching our conclusion, we take into account all the information that is available to us. Suppose, for example, you are trying to decide whether something somebody has told you is true or not. Your decision will be based partly on your estimate of the likelihood of the event they are describing; but it will also take into account the character of the person who told you, their manner of speaking at the time, their possible motives for lying to you, and so on. This is, then, a quintessentially unencapsulated process, in that there is no limit on the

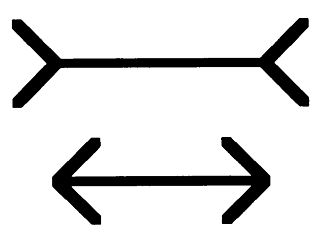

Fig. 9. The Muller-Lyer illusion.

kind of information that might be brought to bear on it. Modular processes, Fodor claims, are encapsulated; central processes are not.

Two points deserve emphasis here. First of all, perceptual representations have semantic content: they are *about* things in the world. We might well think of them as "perceptual hypotheses"—to use Mark DeBellis's useful term (1993, p. 62)—which are sent to the central system for evaluation.[6] But this brings up a second important point: perceptual hypotheses are not in themselves beliefs. If having a perceptual hypothesis P (where P is some kind of proposition) corresponds to a belief at all, it is more like the belief that "my eyes (or ears) are telling me that P." Now, most of the time, we trust our perceptual hypotheses and base our beliefs on them; but when we receive a hypothesis that we know to be false, we are perfectly capable of overruling it. These two points are nicely illustrated by the Muller-Lyer illusion. Here, we are very conscious that the perception we are experiencing is about something; our eyes are telling us quite clearly that the two lines are different lengths. Yet this hypothesis is clearly in conflict with what we believe. As another example, consider the Necker cube, shown in Figure 10. It is possible to see this figure in two distinct ways, either with corner A in front of corner B, or vice versa; one can make one's perceptions of it "flip-flop" back and forth between the two. But as one's perceptions of it change, one's beliefs about the pattern surely are not changing; one continues to know that it is a single pattern that can be seen in two different ways. What is changing, according to the modularity view, is our perceptual hypotheses. One of the great attractions of the modularity theory,

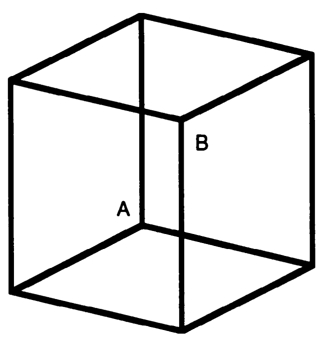

Fig. 10. The Necker cube.

6. The term "hypothesis" is perhaps not ideal here. Having a perceptual hypothesis is clearly different from entertaining a hypothesis about something hypothetical. We can imagine all kinds of hypothetical propositions that conflict with our beliefs: for example, that dogs were green or that the world was flat. But clearly, this is qualitatively different from experiencing an incorrect perceptual hypothesis: for example, the hypothesis in the Muller-Lyer illusion that the two lines differ in length.

in my view, is its ability to account for mysterious phenomena such as these.

If we accept that there are some perceptual hypotheses that are fast, mandatory, and encapsulated, this raises the question, What kinds of things can perceptual hypotheses be about? In Fodor's words, the outputs of modules are "shallow." This entails, first of all, that modular judgments are concerned only with things in the immediate visual or auditory scene: the phonological and syntactic structure of a linguistic utterance or the spatial arrangement and nature of objects in our visual field. Beyond this, it must be admitted that Fodor is rather unclear, and indeed rather inconsistent, about what kinds of judgments modules perform. In *The Modularity of Mind,* he proposes that vision modules might perform certain kinds of basic object identification—for example, identifying things as dogs or chairs, but not as silver-haired poodles or wing-back armchairs; elsewhere, however, he seems to suggest that even basic kinds of object identification are central processes rather than modular processes.[7] A further point seems clear, however: because modules are informationally encapsulated, there must be limits on the kinds of information that can be brought to bear on modular judgments, even regarding the immediate environment. For example, we might suppose that our vision module was capable of identifying dogs; but only if we can claim that such judgments are based only on certain information (presumably information such as the spatial shape or movement of the object) and are not influenced without limit by other information we might possess or obtain: for example, being told that something was a dog, or inferring that it was a dog because we know dogs are common in the neighborhood. This may seem to render Fodor's thesis somewhat implausible (at least as it pertains to dog identification); but it seems to me necessary in order to preserve encapsulation. If *any* information we possessed could influence the module's decision of whether something was a dog, then encapsulation would have no meaning.[8]

7. The view that object identification is modular is presented in *The Modularity of Mind,* pp. 94–97; for the view that it is a central process, which merely takes the module's output as a starting point, see "A Reply to Churchland," p. 259.

8. A good deal of recent experimental work on modularity has focused on this issue, particularly with regard to language. Fodor has hypothesized that, in language perception, words are identified, and a syntactic structure is generated, in a modular way, without influence from higher knowledge; this modular representation is then passed on to the central system for higher level processing and evaluation of meaning. But if the lower levels of language perception are indeed modular, it should be the case that higher knowledge about the context and situation (i.e., our expectations about what the speaker is *likely* to say) has no effect on these lower levels. For example, knowing what the speaker is likely to say should not affect the speed of syntactic processing or the choice (at least, the *initial* choice) of syntactic structures from among the possible alternatives. A number of experimental studies have been done on exactly this question; so far, the evidence is mixed. A recent anthology, *Modularity in Knowledge Representation and Natural-Language Understanding* (Garfield, 1987), contains several studies dealing with this issue.

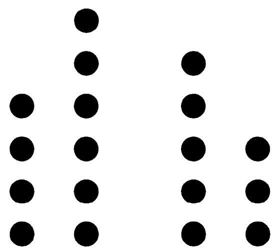

Fig. 11.

It is equally important to stress, however—and this is something that Fodor does not discuss—that even some judgments that are clearly about the scene before us, and that rely on very limited information, appear to be nonmodular. Consider Figure 11. Suppose you were asked whether these two pairs of columns, when read as two two-digit numbers, added up to 100. This is a judgment that is clearly about the immediate scene; moreover, it is one that does not depend on outside information in any obvious way (the answer is going to be the same regardless of external circumstances); in this sense it could, on principle, be made within the module. But would this judgment be modular? It would not be fast; it would not be mandatory. More importantly, it would not be accompanied by any experience such that we would say "these pairs of columns *look* like they spell two-digit numbers that sum to 100" (as opposed to, for example, seeing two parallel lines that look like they are the same length). This experiential aspect is important, and I will return to it later. In any case, the point for the moment is that even some judgments that are quite clearly about the scene before us, and highly limited in the information they draw on, may still be nonmodular.

It is important to stress that Fodor readily concedes that the capacities of one's modules may change over time. He maintains, for example, that lower level language processing (such as the identification of words) is a modular process, and this is obviously something that is learned. Another example is the inverted-lens experiment. When people are fitted with special goggles that make everything look upside-down, after a few weeks they adjust and begin to see everything right-side-up again. The important point for Fodor, though, is that modules learn on their own, not by influence from central processes. After all, with the inverted-lens experiment, people know immediately that everything is upside-down, but it still takes their perceptions several weeks to adjust to the fact. Thus Fodor might readily concede that, if we practiced enough, we could come to identify patterns as having columns of dots that spell two-digit numbers summing to 100, or as having any other property, in a modular fashion; the important point is that we do not perceive them that way now.

In a sense, of course, it is nothing new to claim that perceptions represent only part of our knowledge about a situation. For there is no doubt that our perceptual experiences of things—our experiences that things look or sound a certain way—are unaffected by all kinds of changes in our knowledge about them: a dog *looks* exactly the same regardless of what its name is, for example. But one could concede that perceptions represent certain kinds of information and not others without accepting that perceptions were in any way "encapsulated." The truly convincing evidence for Fodor's claim, I think, is things like the Muller-Lyer illusion. The information involved here—the relative length of parallel lines—is clearly a *sort* of information that perceptual representations do contain. We can feel our eyes telling us that the lines are different lengths: yet we know they are the same. How can one explain this phenomenon without allowing that perception is (at least sometimes) encapsulated from belief? It is admittedly odd, however, that so few clear examples of encapsulation can be found. One might accept things like the Muller-Lyer illusion as pointing toward encapsulation. But if such artificial examples were the only cases that could be found, the remarkable thing would surely be how rarely our perceptions mislead us—or even fail to tell us what we want to know. Indeed, Fodor's main aim in *Modularity of Mind* is to suggest that something like encapsulation might exist and to shed doubt on some evidence that was being put forth as proving plasticity (pp. 73–86). As he admits, the positive evidence he presents for encapsulation is rather meager.

Modularity and Motivic Perception

The possibility that musical perception might be modular has been explored by several theorists. In his book *Consciousness and the Computational Mind* (1987), Ray Jackendoff incorporates his and Lerdahl's theory of musical perception into a broad theory of perception and cognition that draws heavily on Fodor's ideas.[9] Eugene Narmour (1990) also invokes the idea of modularity in his theory of melodic perception. Other theorists, notably Isabelle Peretz and Jose Morais (1989), Naomi Cumming (1993), and Mark DeBellis (1993), have explored the implications of modularity for musical perception in a general way. To my knowledge, however, the possibility that motivic structure might be perceived in a modular fashion has not been addressed. Let us examine this possibility.

The modularity theory has some immediate attractions as an explanation of motivic perception. Motivic perceptions are shallow, in that they are about the auditory scene before us. Some motivic perceptions are surely fast and mandatory. Hearing the relationship in the Mozart is instanta-

9. See also Jackendoff's article "Musical Processing and Musical Affect" (1991).

neous and automatic; we hear it whether we want to or not.[10] Other motivic relationships, however, are not perceived in a fast mandatory manner, such as the retrograde in the Webern. Such relationships can be detected through listening, but it takes time and deliberate effort. Or consider the two settings of the Haydn melody. In Figure 5, the relationship may be detected, but only with effort; in Figure 6, when reinforced by the meter, it is detected automatically. Purely from a computational point of view, then, this raises the possibility that some relationships are perceived in a modular way but others only in a "central" way.

However, there is a phenomenological side to this as well. Consider the Mozart symphony melody. It seems to me that in this melody, there is a strong sense that our experience of the passage has propositional content: that is, we have a sense of hearing *that* the two segments are related, just as there is a sense of, for example, seeing *that* two lines are the same length. (It seems to me that when people speak of "hearing a relationship" this is usually the sense they intend. Henceforth I will assume this usage: to hear a relationship means to hear that it is there.[11]) In the Webern, however, there is no such sensation of hearing the retrograde in this sense, even if we know it is there. The case is perhaps clearer in the two Haydn examples. In Figure 6, our experience seems to contain the information, "there are two parallel fragments there." Figure 5, however, is qualitatively different: there is no sense of hearing that the relationship is there, even when we are actually detecting it. This suggests that hearing a motivic relationship is not simply the same as detecting it; the latter can occur without the former. Now, there is a possible objection to this argument. One might say, in the case of the Haydn, "but in hearing Figure 5 and detecting the relationship, we are in fact hearing that it is there just like we are with Figure 6. This is what it is

10. Even in pieces where there is otherwise very little sense of meter or motivic structure—and where listeners are thus not expecting to hear structures of this kind—any passage that does suggest a metrically parallel motivic pattern will tend to stick out very strongly, as composers of such music are well aware. The fact that such patterns are heard in the absence of any expectation of them (let alone any conscious strategy of *trying* to hear them) is further evidence for the "mandatoriness" of motivic perception.

11. Of course, "hear something" has a more general meaning as well; "to hear X" can mean something like "to aurally receive an acoustic signal that contains X," even if that information is not in any way extracted. In this more general sense, we are clearly "hearing" the retrograde in the Webern simply because we are receiving the signal that contains it. That is, of course, not the sense of "hear" that I intend here. A further point: "to hear that" has another meaning as well, meaning to "come to know through auditory perception, directly or indirectly". In this sense, we might "hear that" Joe had arrived by being told about it, hearing shouts of joy, etc. This is again not the sense of "hear that" that I intend; the sense I intend entails "to come to know something aurally," but with the added clause that the information is somehow *part* of the experience, rather than something inferred from the experience.

like to hear that a nonmetrically parallel transposition is there; hearing nonparallel transpositions just feels different from hearing parallel ones." To detect a relationship is to hear that it is there, one might say, no matter what kind of relationship it is; it's just that different kinds of relationships sound different. But if this is true—if our hearing of Figure 5 contains the information that the relationship is there—then our experience of the passage should differ depending on whether we detect the relationship or not. It seems to me that it does not; detecting the relationship has no effect on the *sound* of the passage. Similarly, our experience of the Webern passage is unaffected by our detection of the retrograde: It sounds the same in either case.

Phenomenology suggests, then, that there is a process relating to the identification of motivic relationships that is distinct from the simple detection of those relationships. In some cases, we have a strong sense of hearing that a certain relationship is there (Figure 6); in other cases (Figure 5), there is no such sense, although we may be detecting the relationship (and although the relationships are quite similar in the two cases). Another way of looking at this is that it represents a disparity between perception and belief. In some cases when we detect relationships, we seem to hear them as well; in other cases we do not. Again, there is nothing strange about the fact that certain kinds of information are represented in perception and others are not; but the fact that certain motivic relationships are represented perceptually, and other (ostensibly quite similar) ones are not, seems rather mysterious.

By a modularity view, however, these phenomena are accounted for quite nicely. Let us suppose that our hearing of motivic relationships represents the output of some kind of motivic module that is able to detect certain kinds of relationships but not others; others may be detected, but only in a nonmodular, "central" manner. By this view, the reason our hearing of the Haydn differs between Figure 5 and Figure 6, even if we are detecting the relationship in both cases, is that in one case the pattern is being detected by the module, while in the other case it is not. The reason Figure 6 sounds the same whether or not we detect the relationship is that when the relationship is detected, it is being done in a central, nonmodular manner, and this has no effect on the module. This view suggests that there is an important distinction to be drawn between those relationships that are perceived in a modular way—we might call such relationships "phenomenologically direct"—and those that are not. It is interesting to note also that the relationships that are phenomenologically direct—namely, metrically parallel transpositions—are exactly those that are perceived in a fast, mandatory way. The fact—if it is indeed generally true—that those relationships perceived in a fast and mandatory way are experientially distinct from others as well would seem to yield an important confirmation of Fodor's theory.[12]

I have argued that in both the Haydn Figure 5 and the Webern, the relationships in question are detectable only in a nonmodular way; detection is slow and deliberate, and in neither case is the relationship really heard even when it is being detected. However, there is also a subtle difference between the Haydn and the Webern. In the case of the Webern, I do not want to say that my ears are telling me that there is *no* retrograde here. Rather, my ears are expressing no opinion on the subject; they have no concept of retrogrades. In the case of the Haydn melody in Figure 5, however, I feel my ears are telling me quite positively that there is *no* transpositional relationship between the first three notes and the second three. This is, I submit, a form of illusion: a direct conflict between perception and belief.[13] Further evidence of this is the fact that it is rather surprising to learn of a relationship such as that in the Haydn (particularly given its recurrences later in the phrase); and why would we be surprised to learn about it unless we had been perceiving that no such relationship was there? If I am correct, then the Webern example is like the dots in Figure 11; our music motivic simply has no concept of retrogrades, any more than our vision modules have a concept of dots spelling numbers that sum to one hundred. Figure 5, on the other hand, is more like the Muller-Lyer illusion: it is a case of actual conflict between perception and belief. But whether Figure 5 is an actual case of illusion is, in a way, secondary; the main point about this example is that, as in the Webern, the motivic relationship at issue is not being recognized in a modular way.

I have suggested that these cases are not unlike certain phenomena in visual perception; but there is an important practical difference. In vision, our powers of modular pattern recognition seem so great that they tell us

12. It is important to note that saying that some relationships are detected quickly and automatically, while others are not, is quite different from saying that there is an experiential difference between the two kinds. (However it feels to detect a relationship in a nonmodular way—as in Figure 5—one could imagine this experience happening in a fast, mandatory way). The fact that motivic relationships that are detected in a fast, mandatory way are also experientially different from those that are not is an empirical observation, one that seems to accord nicely with Fodor's theory.

13. Whether this is a real case of conflict between perception and belief depends on the kind of perceptual hypothesis we experience when we hear Figure 5, knowing that the parallelism is there. If the perceptual hypothesis is "There is no transpositional relationship between the first and second three notes at all," then the hypothesis is false, we know it is false, and encapsulation is beyond dispute. However, one might say that actually, the perceptual hypothesis we experience is "There is no *metrically parallel* transpositional relationship between the first and second three notes"; in this case, the hypothesis is true. By this view, of course, there should be no surprise upon discovering a relationship such as that in Figure 5; there is nothing to be surprised about. Which of these two views is correct is, I think, a difficult question. But the main point is that the perceptual hypothesis is *not* telling us that the two segments *are* related; this, I feel, is fairly clear.

practically everything we want to know; the demonstrations of modularity discussed in the case of vision therefore seem rather artificial. The fact that we fail to perceive the relationship between the groups of dots in Figure 11, or that our perceptions misinform us about the Muller-Lyer illusion, seems of little practical consequence. In music, however, I shall argue that the distinction between modular and nonmodular motivic relationships forms an important dividing line among the kinds of relationships that are discussed in theory and analysis. It may be that many people have an intuitive sense of this division; but again, it is difficult to reconcile with our commonsense notions of knowledge and belief. The modularity theory may therefore have an important role to play in explaining this phenomenon.

Modular Motivic Perception: A Proposal

If we accept that some motivic relationships are perceived in a modular way and others are not, the question becomes: What kinds of motivic relationships does the module recognize? I suggested that it was much easier to hear motivic relationships that were metrically parallel, that is, similarly placed with respect to the meter. When nonmetrically parallel relationships are heard, it is often the case that a competing metrical pattern is being created to make them parallel. I would now like to suggest tentatively that it is *only* metrically parallel segments that are perceived by the module. Let us define this a little more restrictively to say that two segments are metrically parallel only if they are similarly placed with respect to adjacent beats at some metrical level.[14] Thus each segment may be metrically parallel to several other segments, at different metrical levels. Two segments will then be recognized by the module as being motivically parallel if (a) they are metrically parallel and (b) they are identical in rhythm and related by tonal transposition (or exact repetition).[15] Consider Figure 12, the melody from the second movement of Haydn's "Emperor" Quartet. Pairs of segments that would be recognized as parallel by the module are marked above the staff (not all pairs are shown); some pairs that would *not* be recognized as parallel are marked below. Segments A and B, for example, are metrically parallel, because their beginnings are a half-note apart, and they are related by transposition. Thus they are recognized by the module as motivically parallel. Segments B and C are also motivically parallel, at the whole-note level. The second four-measure phrase of the melody—an exact repeat of

14. Another way of saying this is that two segments are metrically parallel if their beginnings are separated by a unit of the meter, where a unit of the meter is the distance between adjacent beats at some metrical level.

15. Actually, we could consider exact repetition to be a case of transposition: in set-theoretical terms, it is transposition where T=0.

Fig. 12. Haydn, String Quartet op. 76 no. 3 (the "Emperor"), II.

the first—is parallel to the first phrase at the four-measure level of the meter. Notice, however, the segments that are related by transposition, but are not parallel, such as segment A and segment E, or segment B and segment D. These segments are not only related by transposition, but are at the same pitch level as well (that is, transposition is T_0). Yet they are not metrically parallel and thus are not recognized in a modular fashion. I feel that this is phenomenologically exactly right. The metrically parallel segments are exactly those that leap out at us as being parallel; they are the ones that *sound* parallel. Consider another example, the aria "Se vuol ballare" by Mozart, shown in Figure 13. There are again a number of short parallel fragments here, such as segments A and B. In terms of rhythm and intervallic pattern, segments C and D are parallel, as are segments E and F. But these are not metrically parallel and are thus not perceived in a modular fashion. Here, too, it seems to me that it is precisely those fragments that are metrically parallel that sound similar. Again, the point is not just that some fragments are identified more quickly and automatically than others, but that even when one does detect nonmodular relationships, they seem qualitatively different from those identified in a modular way.

Several points should be made about this system. It may be noticed that the parallel segments shown in Figures 12 and 13 constitute, in effect, a parsimonious representation of the corresponding melodies. When a parallelism is found, only the pitch information of the left side of the relation-

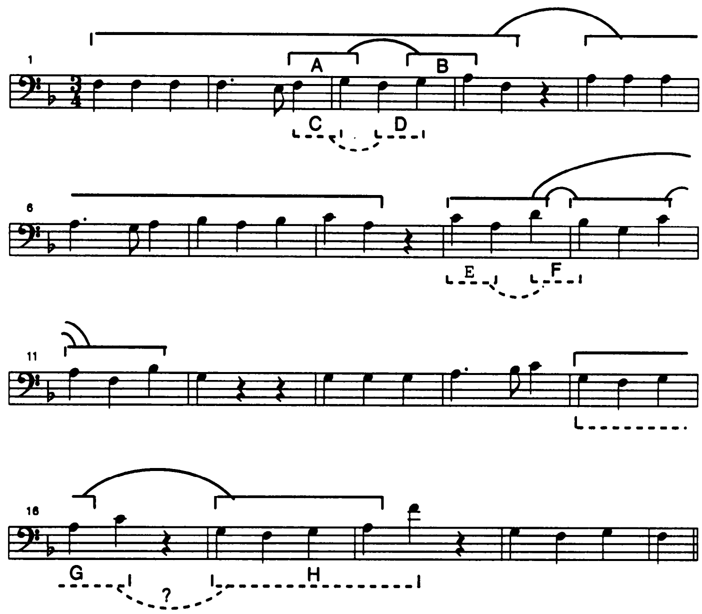

Fig. 13. Mozart, "Se vuol ballare," from *Le nozze de Figaro.*

ship need be stored; the right side is simply represented as being a transposed repetition of the left side. It is surely plausible, for example, that a parallelism such as that between the second four measures and the first four measures of the Mozart would make the melody easier to remember. Although these representations are in a sense parsimonious, they are also very redundant, in the sense that a single event may be encoded several times, as part of different segments at different levels or overlapping segments at the same level. For example, in the third full measure of the Haydn, the first note C is part of segment D, parallel with segment C; but it is also part of segment F, parallel with segment G.[16] But redundancy does not argue against this system as a model of perception: We know that redundancy operates in many kinds of perceptual and information-processing systems, such as language perception.[17] Another possible criticism of the

16. Notice that I do not allow the left side of a pair to overlap with the right side. There is also great redundancy between segments at different levels. For example, in the Haydn, segment I is parallel to segment H at the half-note level, but is also part of a larger four-measure parallelism with the first phrase.

17. As Michael Garman points out, the evidence of redundancy in language perception is all around us: consider that we are able to converse in all kinds of less-than-ideal situations, with background noise, competing conversations, and so on, without making special adjustments either as speakers or listeners. This suggests that in normal conditions we are getting much more information than we actually need. This has also been experimentally verified, for example, in experiments in which phonological segments are systematically deleted from recordings of speech without affecting comprehension (Garman, 1990, pp. 185–191).

system is that there are other kinds of motivic relationships that are recognized by the module and should be incorporated into the theory. I am thinking particularly of pairs of segments that are alike in contour but not exact intervallic pattern, such as segments G and H in the Mozart. It seems to me that such relationships are recognized in a spontaneous, immediate way and are experientially similar to exact transpositions; perhaps, then, they should be considered modular as well. Of course, allowing such relationships would substantially increase the number of relationships posited by the theory. But because only metrically parallel segments would be recognized, the number of relationships found would still be confined to a tiny fraction of those that might be posited in analysis.[18]

If we accept that only motivic relationships between metrically parallel segments are recognized in a modular way, this implies an extremely close connection between metrical structure and motivic structure.[19] The metrical structure, in effect, selects certain segments as being metrically parallel to one another; from among those, the motivic structure finds pairs of segments that are motivically parallel as well. So far, we have assumed that the metrical structure for a piece is already established when the motivic search begins. In fact, this is not the case at all. As Lerdahl and Jackendoff point out, parallelism is itself one of the major factors that determine metrical structure. This can be seen in Mozart's 40th, for example; surely the parallelisms in this melody are, in part, what clue us in to the metrical structure. Recall also the point that was made earlier. Sometimes, even when

18. Recent work in contour theory may be of relevance here; for a survey of this work, see Morris (1993).

Several other possible criticisms of the theory should be mentioned. One concerns tonality. Lerdahl and Jackendoff, in their brief remarks on associational structure, suggest that the associations we hear are affected by reductional structure: that is, by factors of pitch stability, prolongation, and the like (1983, pp. 286–287). It seems to me that this is not really true. Consider the passage shown in Figure 14. Here there is a clear structural difference between the first four notes of the melody and the second four. In the first group, the Cs are harmonically much more stable than the Bs, because they are consonant with the bass; in the second pair, however, the Gs are much more stable than the As. Surely *GTTM* would assign different reductional structures to these groups. Yet to me there is a very strong associational link between the two segments. I suggest, then, that motivic structure is essentially independent of reductional structure.

Motivic structure is clearly not unaffected by tonality, however; the present theory assumes that only tonal transpositions, rather than exact transpositions, are recognized by the module. How the module would deal with chromatic alterations, or with melodies that are not presented in a tonal framework, is an important question that remains to be resolved.

Finally, it may be noted that the current system regards all motivic relationships as pairs of segments. This may seem odd; in the Haydn, for example, rather than hearing several different pairs of descending two-note segments, it might be argued that we hear a single two-note motive repeated several times. This aspect of the theory is admittedly rather counterintuitive.

19. From now on, I will use "motivic structure" to refer only to the structure that I claim is formed in a modular way. It is important to remember, however, that other motivic relationships may be recognized in a nonmodular way.

Fig. 14.

a metrical structure has been firmly established, a motivic pattern is strong enough to be perceived even when it conflicts with the metrical structure; but in such cases, one feels that a competing metrical structure is also being heard. I propose the following explanation for this. The only motivic relationships that are heard are those that are parallel in the current metrical structure. Most of the time, only one metrical structure is being considered. However, we are also continuously reevaluating our choice and weakly considering other metrical structures. One of the factors we consider is the strength of the resulting motivic structure. (There are also other factors affecting metrical structure, unrelated to motivic structure; these are precisely the other metrical preference rules proposed by Lerdahl and Jackendoff.) If we find a metrical structure with a strong motivic structure, that is, one with many parallelisms, that metrical structure will be preferred (particularly if the motivic structure for the current meter is weak, as in the case of the Beethoven passage in Figure 7). It is possible, then, that during points of rhythmic instability, several metrical structures may be present and several competing motivic structures may be heard; although most of the time, it seems to me, a single metrical structure and motivic structure is strongly preferred over all others. In any case, the claim still stands that wherever one hears a motivic relationship, one should hear a metrical structure that goes along with it. Thus we may continue to posit an extremely intimate link between metrical structure and motivic structure. But the relationship between them is two-way. The metrical structure determines the motivic associations we hear, but the strength of associations that are formed in turn affects the strength of the corresponding metrical structure.

Further Implications

In a sense, then, the generation of the metrical structure and the motivic structure for a piece may be seen as a single complex process. If motivic structure is formed in modular fashion, it seems clear that the formation of metrical structure must also be a modular process rather than a central one, because the essence of modularity is that central processes may not affect modular ones. Indeed, particularly in view of the interaction between them, it seems reasonable to regard the motivic and metrical structures as both being formed within a single module. On the whole, this seems plausible: metrical structures, like motivic structure, are formed in a fast, mandatory way. We should note also that according to Lerdahl and Jackendoff's theory, there are strong rules governing the kinds of metrical structures that are well-formed. Well-formed metrical structures must have evenly spaced beats at multiple levels, with a strong beat at any level also being a beat at every lower level. Thus the structures in Figure 15a are well-formed; those in Figure 15b are ill-formed.[20] These constraints add some empirical substance to the claim that only well-formed metrical structures are formed in a fast, mandatory way. Notice also that this further constrains the possible motivic structures: One could imagine motivic structures that were parallel according to the meter in Figure 15b, but because this metrical structure cannot arise, the associated motivic structure cannot arise either.

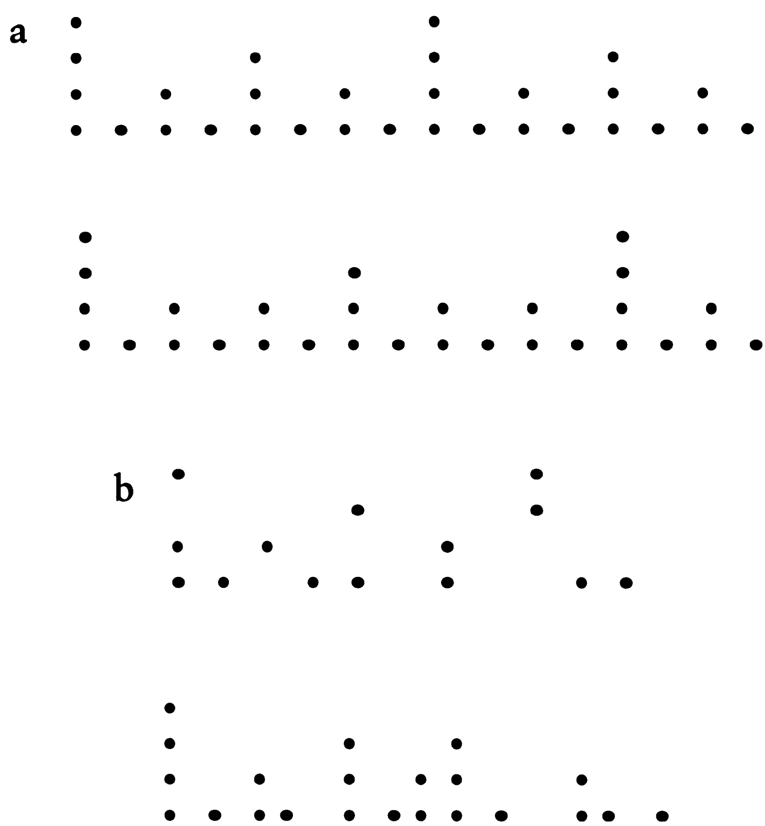

Fig. 15. (a) Well-formed metrical structures. (b) Ill-formed metrical structures.

20. See Lerdahl and Jackendoff (1983), pp. 68–74.

Thus the possible motivic and metrical structures that may be formed in a fast, mandatory way are quite tightly constrained.

We should note, however, that to argue that motivic-metrical structure is mandatory is somewhat problematic. It does seem clear that there is a mandatory link between motivic structure and metrical structure: we cannot hear a motivic structure in the manner just described (fast, phenomenologically direct) without also hearing the associated metrical structure. Yet it is also clear that the metrical structure we impose on a piece—and, by extension, the motivic structure—is to some extent a matter of choice. If one tries hard enough, one could, for example, hear the Mozart aria in $\frac{4}{4}$ or the Haydn "Emperor" melody in $\frac{3}{4}$. And, by choosing these metrical structures, one could indirectly choose the motivic relationships one hears. How can this be reconciled with the presumed mandatoriness of metrical and motivic structure? The mandatoriness issue is a difficult one, which I have to some extent evaded up to now. I have claimed that certain relationships are heard without effort, whereas others (e.g., retrogrades) can only be heard with effort. But in fact, to say that something is heard without effort is not the same as saying it is mandatory: if it may be overruled with effort, then it is not mandatory. (On the other hand, any kind of process that occurs only with effort is clearly not mandatory: effortlessness is necessary, but not sufficient, for mandatoriness.)

It is interesting to note that Fodor has discussed an analogous problem in the debate over modularity.[21] Fodor claims that the operation of the vision module is mandatory; we have no choice as to how it analyzes a visual scene. But, as the philosopher Paul Churchland points out, in the Necker cube, it is possible for us to flip back and forth between the two interpretations at will; how can this be reconciled with mandatoriness? Fodor's response is as follows. In perceiving the Necker cube, he suggests, we do not choose to see the pattern with corner A in front, the way we choose to do an arithmetic problem. Rather, we alter our fixation point; this in turn makes our vision module see the pattern with corner A in front. To say that seeing the pattern in a certain way is a controlled by choice is like saying that our heartbeat is controlled by choice, because we can choose when to take a nap. We can affect the behavior of our modules, Fodor argues, but only indirectly. Fodor's response is clever, but, to my mind, not totally convincing. Even if we conceive of Necker-cube alteration as an indirect effect of central processes on the module, it is an effect nonetheless, and it would seem to endanger the mandatoriness claim. Notice, moreover, that once we lose mandatoriness, we lose encapsulation as well. The whole point of encapsulation is that certain perceptual hypotheses will be

21. Churchland's comments are found in "Perceptual Plasticity and Theoretical Neutrality: A Reply to Jerry Fodor," p. 260. Fodor's response is in "A Reply to Churchland," pp. 255–257.

formed regardless of information that we may possess centrally, for example, knowledge that the hypotheses are false. But if a modular process is not mandatory, then central processes can modify it or alter it; presumably this can be done on the basis of any information the central processes happen to have. In that case, in what sense is the process "encapsulated"?

The response I propose to this objection is somewhat more compromising than Fodor's, but also, I think, more convincing. I have argued that certain metrical-motivic relationships are perceived in a certain way—quickly and effortlessly—and that these structures are subject to strong constraints, regarding the relationship between the motivic and metrical structures and the nature of the metrical structure. We may then concede that these structures are not mandatory and may be overruled; but when they are overruled, they can be replaced only by another structure that is subject to the same constraints. That is, we can choose to hear Figure 12 with a $\frac{3}{4}$ metrical structure and the associated motivic structure. But we cannot choose to hear it with a poorly formed metrical structure such as those shown in Figure 15b, we cannot choose to hear it with a $\frac{3}{4}$ metrical structure and a $\frac{4}{4}$ motivic structure, and we cannot hear it with no metrical-motivic structure at all. In allowing that we can choose to suspend one motivic-metrical structure and initiate another, we must allow that the replacing structure is heard with effort and therefore is clearly not "effortless." But these structures are phenomenologically direct and subject to the same structural constraints as the effortless ones; therefore they seem best regarded as "modular" also. Furthermore, there are limits on how these structures (both the effortless and replacement ones) can be used; one can be replaced with another, but they cannot be suspended entirely, nor replaced with structures of a different kind.

Because the phenomenologically direct structures that are formed are so constrained, and because our intervention with these structures seems to be limited to replacing one structure with another one of similar character, it seems most plausible to view the situation essentially as Fodor describes it. In overruling a modular metrical-motivic structure, our central system is sending a message to the metrical-motivic module, telling it to choose another structure. (There is an intuitive appeal to this explanation; in forcing myself to hear a wrong interpretation of the Haydn melody, I feel I am sending an internal "beat" to the module that confuses it into imposing a different structure.) Our influence on the metrical-motivic module is limited to this. Contrast this with the case of phenomenologically indirect relationships, where we are have tremendous freedom as to what relationships we detect and when we detect them. Therefore, although we cannot say that modular relationships are either mandatory or effortless, we might call them "automatic": generally effortless and affected by choice only within narrow limits. As such, they are clearly set apart from nonmodular relationships, which are always effortful and entirely controllable by choice.

What can be said about the other characteristics of modular outputs? According to Fodor, the outputs of modules are "shallow": they are about the immediate scene before us. This distinction is of little use to us here, as presumably all motivic relationships are shallow in this sense. Regarding informational encapsulation, the situation is more complicated. Fodor's prime cases of encapsulation are cases of illusion: cases in which our perceptual hypotheses appear to actually conflict with our beliefs. I argued earlier that Figure 5 might represent such a case; however, I do not wish to insist on this argument, and I feel that the case for modularity is convincing without it. I also argued that, for judgments to be modular, there must be limits on the information that can influence them. Nonmodular judgments, as defined here, are virtually unconstrained: we may detect all kinds of relationships, even those that might only be learned about through explicit theoretical training (I am thinking here of Z-related sets and other relationships discussed in set theory). Metrically parallel transpositions, however, appear to be perceived even by untrained listeners, without any theoretical guidance; it seems plausible to believe, then, that this perceptual capacity develops entirely within the module, without any guidance from central processes. This is another sense, then, in which the modular relationships discussed here appear to be encapsulated.

An important point should be made about the claim that the motivic capacities of the module are not influenced by central processes. This is not to say that the module's abilities are innate: they might be learned over time, as long as this development takes place within the module. This is exactly parallel to the case of language in Fodor's theory. Nor—even more importantly—is it to say that people could not learn to perceive other kinds of relationships in a modular way. Fodor allows that modules can learn to perceive things, and it is conceivable to me that with enough exposure, one's module might learn to perceive nonmetrically parallel transpositions, retrogrades, and all kinds of other relationships. My only claim is that at the moment, I (and I believe most other people) do not perceive them in a modular way. And again, the important point is that if our modules do learn to hear new things, they do so without guidance from our theoretical knowledge and conscious goals; rather, they learn on their own (through a process that Fodor has little to say about).

The modular view of motivic structure offers some insight into the question of musical value and enjoyment. The motivic structure of the Haydn "Emperor" melody, for example, is a complex one, involving parallel relationships at several metrical levels, of different lengths and at different locations. Other things being equal, we may continue to search for parallelisms that have the same metric placement as those found previously; but failing that, we may have to compare different segments, placed differently relative to the beat or at other metric levels, or even segments that are not parallel at all given the current meter (thus requiring a different metrical structure). In this respect, the process of listening to a melody is a process of searching for order and pattern. In this sense, it perhaps has something in common with central processes such as solving a detective or mystery story, solving a riddle, or playing games such as charades or 20 questions; in each case, part of the appeal is the process of finding the order or pattern in a seemingly chaotic or disorderly situation. The idea that the listening process is an attempt to organize what one hears is of course not new; it is reflected most clearly, perhaps, in the writings of Leonard Meyer.[22] One appealing aspect of the current theory is that it offers a way of integrating metrical structure into this view. If one believes that listening to music is a process of finding order, one must ask, what role does meter play? In what sense are we searching for order when we impose a metrical structure on a piece? The current theory offers an answer: Imposing a metrical structure on a piece is nothing but comparing certain segments in search of a parsimonious encoding of the music. In imposing a certain metrical structure, we are essentially guessing that certain segments will be similar to certain other segments in intervallic pattern; finding these patterns will allow us to encode the piece more efficiently.[23]

There is a problem, however, with this view of musical perception. If the process of listening to music is a process of searching for order or pattern,

22. See especially *Emotion and Meaning in Music* (1956), pp. 83–91, and *Explaining Music* (1973), pp. 3–5.

23. I stated earlier that the current theory applies only to low-level motivic relationships. The reason for this lies in the connection with meter. For two segments to be motivically related, they must be similarly placed with respect to adjacent beats at some metrical level. In Lerdahl and Jackendoff's view, however (and I agree), metrical structure only extends up to a certain point. That is, given a piece that falls into 4-measure phrases, we may hear the downbeat of each fourth measure as strong, but we may go no further; rather than hearing every second or third of *those* beats as stronger than the others, we may hear them all as essentially equal. The current theory actually sheds some light on this. Let us suppose that imposing a metrical structure is nothing but comparing two segments in search of a parsimonious encoding of them. That is, if we impose a metrical level at the 4-measure level (so that every fourth downbeat is strong), this is to say that we compare each 4-measure segment with the following one, hoping that the two segments (or parts of them) will be motivically related. In this sense, motivic structure, and hence metrical structure, clearly depends on memory: in order to compare two segments, we must memorize the earlier one. At a certain level, however, the length of the segments involved exceeds our short-term memory capacity. Therefore we might be able to compare 4-measure segments, but not, say, 8- or 12-measure segments, because those would exceed our memory. At this level, then, the segment-comparison process could not be executed, and the sensation of meter would correspondingly fade.

A word is needed here about long-term motivic relationships. It seems to me that these are qualitatively different from short-term relationships. Long-term relationships are usually heard only when the original motive has been repeated several times and thus has been committed to long-term memory. In that case, the identity of the motive is abstracted: it is still remembered as occurring at a certain position relative to the metrical structure, but it will now be recognized whenever it occurs in that position, regardless of its exact distance from the original appearances of it. But this all has to be worked out more fully.

and if musical enjoyment results, in part, from this process, then it is difficult to explain why we could continue to enjoy listening to pieces that we know well, or even know from memory. If we already know a piece from memory, presumably we have already discovered whatever order there is to be found; we already have it parsimoniously encoded in some way, and there is no need to repeat the process. Indeed, once we know a piece well, it is hard to see how the process of hearing it would have any effect on us at all. As Jackendoff has pointed out, however (1991, pp. 226–228), it is precisely here that the modularity view of music yields perhaps its greatest payoff. Let us suppose that the process of analyzing the melody, described above, is performed by a module, or an automatic machine with no access to our background knowledge. The machine does not know we have heard the melody before, and it has no access to any representation of the piece that we might store elsewhere in memory. It simply processes the piece as if for the first time, every time it hears it. This would perhaps allow us to reconcile the fact that the process of listening seems to be a process of learning or discovering with the fact that we seem to be able to derive enjoyment from this process over and over again. In this sense, being a modular process, listening to music would not be so much like a central process such as solving detective stories or intellectual puzzles; rather, it would be more like other modular processes such as parsing sentences or analyzing visual scenes. This then raises another question: why is it that parsing melodies is, in itself, a source of pleasure, whereas parsing sentences and visual scenes generally is not? I have no answer for this; I can only suggest that, perhaps, a modularity view takes us from a baffling situation to a merely highly mysterious one.

Let us summarize the conclusions of this paper. (1) Certain motivic relationships are perceived in a fast, automatic way (automatic in that it is often effortless, and controlled only to a limited extent by choice). (2) Only certain relationships are perceived in this way: metrically parallel transpositions. (3) Such relationships are also phenomenologically direct: there is a strong sense of hearing *that* they are there. (4) Other relationships—nonparallel transpositions, retrogrades, and other set-theoretic relationships—may be detected, but only in a slow, deliberate, phenomenologically indirect way. (5) Metrical structure, too, is perceived in a modular—fast, automatic—way. (6) Metrical structure does not merely determine motivic structure, however; it is also greatly affected by it, in that we prefer to form a metrical structure that results in a strong motivic structure.

In proposing that musical perception is modular, I do not wish to downplay the importance of central processes in musical experience. With regard to the classical-period pieces discussed here, for example, there can be no doubt that things such as narrative structure play an important role, as well as other kinds of extramusical associations—cultural and personal—

that could not possibly be accounted for as modular phenomena. But to admit the importance of higher knowledge in musical experience is not to deny that some aspects of musical mental representations may be formed independently of this knowledge, in an entirely modular way. I have argued that some aspects of musical experience seem to accord well with, and indeed demand, such an explanation. I have also suggested that musical perception has something to contribute to the debate about modularity, offering, perhaps, a uniquely compelling demonstration of the theory's validity.

References

Browne, R. (1974). Review of *The structure of atonal music* by Allen Forte. *Journal of Music Theory, 18,* 390–415.
Churchland, P. (1979). *Scientific realism and the plasticity of mind.* Cambridge, UK: Cambridge University Press.
Churchland, P. (1989). *A neurocomputational perspective.* Cambridge, MA: MIT Press.
Cumming, N. (1993). Music analysis and the perceiver: A perspective from functionalist philosophy. *Current Musicology, 54,* 38–53.
DeBellis, M. (1993). Is there an observation/theory distinction in music?" *Current Musicology, 55,* 56–87.
Deutsch, D. (1982). The processing of pitch combinations. In D. Deutsch (Ed.), *The psychology of music.* New York: Academic Press.
Dowling, W. J., & Harwood, D. L. (1986). *Music cognition.* Orlando, FL: Academic Press.
Epstein, D. (1980). *Beyond Orpheus.* Cambridge, MA: MIT Press.
Fodor, J. (1983). *The modularity of mind.* Cambridge, MA: MIT Press.
Fodor, J. (1990). *A theory of content and other essays.* Cambridge, MA: MIT Press.
Frisch, W. (1984). *Brahms and the principle of developing variation.* Berkeley, CA: University of California Press.
Garfield, J. L. (Ed.). (1987). *Modularity in knowledge representation and natural-language understanding.* Cambridge, MA: MIT Press.
Garman, M. (1990). *Psycholinguistics.* Cambridge, UK: Cambridge University Press.
Jackendoff, R. (1987). *Consciousness and the computational mind.* Cambridge, MA: MIT Press.
Jackendoff, R. (1991). Musical processing and musical affect. *Music Perception, 9,* 199–230.
Krumhansl, C., Sandell, G., & Sergeant, D. (1987). The perception of tone hierarchies and mirror forms in twelve-tone serial music. *Music Perception, 5,* 31–78.
Lerdahl, F., & Jackendoff, R. (1983). *A generative theory of tonal music.* Cambridge, MA: MIT Press.
Meyer, L. (1956). *Emotion and meaning in music.* Chicago: University of Chicago Press.
Meyer, L. (1967). *Music, the arts, and ideas.* Chicago: University of Chicago Press.
Meyer, L. (1973). *Explaining music.* Berkeley, CA: University of California Press.
Monelle, R. (1992). *Linguistics and semiotics in music.* Switzerland: Harwood Academic Publishers.
Morris, R. (1987). *Composition with pitch-classes: A theory of compositional design.* New Haven: Yale University Press.
Morris, R. (1993). New directions in the theory and analysis of musical contour. *Music Theory Spectrum, 15,* 205–228.
Narmour, E. (1990). *The analysis and cognition of basic melodic structures: The implication/realization model.* Chicago: University of Chicago Press.

Peretz, I., & Morais, J. (1989). Music and modularity. *Contemporary Music Review, 4,* 279–293.
Reti, R. (1951). *The thematic process in music.* New York: The Macmillan Company.
Rosen, C. (1971). *The classical style: Haydn, Mozart, Beethoven.* New York: The Viking Press.
Rothgeb, J. (1983). Thematic content: A Schenkerian view. In D. Beach (Ed.), *Aspects of Schenkerian theory* (pp. 61–76). New Haven: Yale University Press.
Schachter, C. Motive and text in four Schubert songs. In D. Beach (Ed.), *Aspects of Schenkerian theory* (pp. 61–76). New Haven: Yale University Press.
Schoenberg, A. (1950). *Style and idea.* New York: Philosophical Library.
Tovey, D. F. (1939). *Essays in musical analysis* (vols. 1–6). London: Oxford University Press.