Music Cognition, Internalism and the Extended Mind

by

Luke Kersten

A thesis submitted in partial fulfillment of the requirements for the degree of

Master of Arts

Department of Philosophy

University of Alberta

Abstract

This thesis aims to advance the study of music cognition by examining its philosophical foundations. It is argued that music cognition is problematically committed to a particular, questionable philosophical mode of thought about the mind and cognition known as internalism. In response, an alternative approach to music psychology is offered. This wide music psychology presents a non-internalist approach to music cognition research. It is further argued that within wide music psychology music cognition is extended.

Dedication

To my Grandfather

Theo Kersten

who always got me where I needed to go (if not a little early)

Acknowledgments

There are several people I would like to acknowledge for helping to make this thesis possible. I would like to thank my thesis supervisor Rob Wilson for his continued support and assistance during the preparation and writing of this thesis, particularly his lightning-fast editing skills. I would also like to thank Michael Dawson who first turned me onto and got me interested in music cognition as a topic of study. I am also grateful to the graduate students at the University of Alberta for engaging conversations on parts of this thesis, particularly James Bachmann and Emma Chien for discussions at our semi-regular "Stammtisches". Finally, I would like to acknowledge my deep debt to Laura Davis. Without Laura's keen editing powers and patient support this thesis would not be what it is.

Table of Contents

List of Figures

Introduction

In the same year that the Cognitive Science Society was founded, Longuet-Higgens (1979), the man who coined the term “cognitive science”, published a review and synthesis of his research on the perception of tonal relations. This means that as a field of research music cognition has been active almost since the inception of cognitive science. To many, this presence so early in the history of cognitive science signals the importance of music cognition as a topic of study (see Pearce & Rohrmeier, 2012). And yet, in spite of its perceived importance, music cognition has received little philosophical attention, especially when compared to its mother discipline (e.g., Thagard, 2010). This state of affairs suggests that music cognition or music psychology -- I will use the terms interchangeably -- is a rich, untapped resource for philosophical theorizing.

This thesis attempts to begin the mining process. The primary aim is to advance the study of music cognition by examining its philosophical foundations. The thesis looks to examine the philosophical assumptions that have influenced both the theory and practice of music cognition. The central problem this thesis means to address is whether the philosophical foundations of music cognition are explanatorily robust. It answers in the negative. Music cognition problematically remains committed to a particular, questionable philosophical mode of thought about the mind and cognition.

Precisely what mode of thinking needs to change? It is what philosophers of cognitive science and mind have called "individualism" or "internalism" (see, e.g., Burge, 1986; Wilson, 1995).[1] It is a view about how to understand and study cognition, one that, as we will see, makes specific claims about what cognitive states are, how they are to be individuated, and taxonomized. It is, as Wilson says, "a substantive, plausible, and controversial view" (1995, p.1). Historically, the growth of internalism coincides with the growth of music cognition. As a stand-alone field of research, music cognition shares much of its history with the rise of internalism. And, as we will see, there is some conceptual cross-fertilization between internalist approaches to other areas of cognitive science and music cognition.

But what exactly is music cognition? What is its subject matter? For many, music cognition is the study of the processes that underlie the human ability to understand music. At its core, it deals with qualities of sound such as pitch, duration, loudness and timbre and how they are implicated in the perceiving, learning, and remembering of music (Dowling & Harwood, 1986). But also, more generally, it encompasses musical behaviours. It looks to provide an account of how individuals do everything from performing and composing musical pieces to how music affects motor and cognitive development.

[1] In this thesis I take "individualism" and "internalism" to name the same view. I recognize that this has not always been the case. Individualism is usually more strongly associated with what I call the methodological assumption (Burge, 1979), while internalism is more strongly associated with what I call the metaphysical assumption (Menary, 2007). Yet there are times when they overlap (see Wilson, 1995, ch.1).

In this thesis, the scope of what is meant by music cognition is narrowed. Music cognition is taken to refer to the largely subconscious processing of musical information. This refinement is made for two reasons: First, it would be almost impossible to do justice to all of music cognition research if studies where included that dealt with experience and meaning; second, internalism, as a philosophical framework for cognition, has usually been pitched as a view pertaining to processes at the subconscious level. Music cognition's relationship to internalism therefore becomes more conspicuous if we restrict our focus to subconscious processing.

Moreover, while what is meant by music cognition is narrowed, what is meant by "cognition" is broadened. In this thesis, cognition refers to both central cognitive processes and sensory and motor processes. In addition to those processes that allow for capacities such as language, cognition is conceptualized as also including processes that enable perceptual and motor capacities.

What is it that this thesis hopes to accomplish? First, as mentioned, it aims to broach a previously unexplored relationship between internalism and music cognition. I shall argue that music psychology has been committed to internalism and that because of this allegiance it has incurred certain theoretical or explanatory deficiencies (Chapter 1). Second, this thesis aims to provide an alternative approach to music psychology research (Chapter 2). This alternative approach seeks to identify and examine the elements needed to conduct music cognition research on non-internalist or "wide" terms.[2] I call this "wide music psychology". In articulating wide music psychology, the aim is to push research, in some respects, away from its internalist commitments and towards recent externalist theorizing in the philosophy of cognitive science (e.g., Clark, 2008; Wilson, 2014).

On precisely this front, I further argue that within wide music psychology music cognition is "extended". By this I mean that the processes underlying music cognition reach out beyond the boundary of the individual. This seemingly radical conclusion, I argue, is the natural extension (no pun intended) of moving to an externalist-inspired wide music psychology. Finally, the thesis concludes with an evaluation of alternative approaches to cognitive extension within music cognition (Chapter 3). The purpose this is to unify disparate lines of theorizing on cognitive extension and music cognition.

[2] Use of the term "wide" here is meant to contrast the elements of this approach with those of the internalist paradigm.

Chapter 1 - Internalist Music Cognition

1.0 Chapter Overview

This chapter has two goals, namely to: (i) introduce the reader to some important strands to music cognition research, particularly the ways in which it has been committed to internalism, and (ii) put pressure on the internalist dimensions of music cognition. With respect to (ii), I attempt to highlight two explanatory deficiencies for an internalist approach to research. In doing so, I also attempt to motivated an alternative, wide approach to music cognition.

1.1 Internalism and Psychology

The first half of this chapter is devoted to showing that music cognition has been committed to an internalist paradigm of thought. To see how music cognition has allied itself with internalism, the core suppositions of internalism need to first be examined. To achieve this, discussion can be structured around three questions: How should cognition be studied? Where is cognition located? And how does cognition work? Articulating internalism's answer to each question should provide a more robust characterization of what the view amounts to.

1.1.1 The Methodological Assumption

First, according to internalism, when it comes to cognition, factors external to the individual are beside the point. As Burge (1986) characterizes the view:

> According to individualism [internalism] about the mind, the mental natures of all a person's or animal's mental states (and events) are such that there is no necessary or deep individuative relation between the individual's being in states of those kinds and the nature of the individual's physical or social environments. (Burge, 1986, p.3-4)

Cognition is autonomous from what goes on outside the head. What is explanatorily and taxonomically important is that cognitive processes play certain roles within an individual's internal mental economy. Call this internalism's methodological assumption.

The methodological assumption has both a negative and positive form. The negative form is expressed in what's called "methodological solipsism", the claim that "[p]sychological states should be construed without reference anything beyond the boundary of the individual who has those states" (Wilson, 2004, p.77). The positive formulation is expressed in Stephen Stich's principle of autonomy, the view that "[h]istorical and environmental facts will be psychologically relevant only when they influence an organism's current, internal physical state" (1983, p.165).

According to Wilson (1995), the methodological assumption stems from concerns about individuating mental states by causal powers. Since two organisms can have different output with respect to the same input, if psychology is to make projectable generalizations, its individuative practices need to be sensitive to the (internal) causal properties that make the difference. However, this is not to say that the environment is never appealed to in internalist accounts. Rather, it is to say that the internal factors are the primary explanatory focus and historical and environmental factors secondary.

1.1.2 The Metaphysical Assumption

Second, on the internalist view, cognitive states supervene on only the intrinsic, physical properties of individuals. Call this internalism's metaphysical assumption. What does it mean to say that mental states supervene on physical states? It means that mental states of type A depend on or are fixed by physical states of type B such that there can be no difference among A states without a corresponding difference among B states. According to the internalist, the supervenience thesis holds for only the intrinsic physical states of individuals. The types of physical properties that fix mental states are wholly internal to the individual. As Devitt has put it, "[o]nly something that is entirely supervenient on what is inside her skin -- on her intrinsic internal physical state, particular her brain -- could play the required explanatory role between peripheral input and output" (1990, p.377). Two lines of reasoning are generally taken to support the internalist supervenience thesis.

First, there is a "standard intuition" that says that the internalist supervenience thesis should be appealing because it makes intuitive sense of causal behaviour (see Wilson, 1995, p.13). Suppose, for instance, there are two physically identical individuals, what are sometimes called "döppelgangers". If placed in different environments, they may exhibit the same behaviour. Since the environmental properties are different, the only way to account for any similarity in behaviour is in terms of the sameness of their mental states. But if the individuals are constitutionally identical, they have the same intrinsic, physical properties. Thus, the mental states must supervene on the intrinsic physical properties of which they both share. To deny this would be to counterintuitively insist that the two individuals could have different psychological states despite having identical physical states. Intuition seems to support the internalist supervenience thesis.

Second, as the Devitt quotation suggests, there is an explanatory line that says that given that psychological explanations seek to understand how people produce the behaviour they do, the internalist supervenience thesis is the best chance of doing this. Jerry Fodor nicely expresses this position when he writes: “[W]e abandon this principle at our peril; mind/brain supervenience (/identity) is our only plausible account of how mental states could have the causal powers that they do” (1987, p.44). Good psychological explanation entails something like the internalist supervenience thesis. I will return to this second point in section 1.3.1. It is also probably worth mentioning that the internalist supervenience thesis is best understood by its advocates as providing a philosophical gloss on what is implicit in psychological practice.

1.1.3 The Operational Assumption

Finally, for internalism, cognition operates through the manipulation of internal information-bearing structures or vehicles. An information-bearing vehicle or structure is an organization or arrangement of physical properties that allows the carrying of information. For example, tree rings carry information about tree age in virtue of the fact that the physical patterns of the wood co-vary with the age of the tree. Piccinini et al. (2011) call this "natural information". Note that the content (the information or what is represented) is distinguishable from the vehicle (the physical pattern) that carries the information.

In the cognitive sciences, information-bearing vehicles are usually conceived of as internal representations of some kind. Since people can think and act independent of what is present in the local environment, what matters for psychological studies is how people represent the world, not how the world actually is. Cognitive psychology looks to identify the cognitive contributions to behaviour. As Bach says,

> [W]hat is outside a person's mind is irrelevant to psychology. Regardless of how the world is in comparison to how it is represented as being and regardless of how it may change while the person's psychological states remain the same, everything is the same as far as [internalist] psychology is concerned. (Bach, 1983, p.123)

The operational assumption claims that cognition trades in internal information-bearing structures.

1.1.4 Summary

Let us return to the orienting questions with which the section began. First, we asked how cognition should be studied. The response from internalism was that cognitive states are to be taxonomized and individuated by reference to intrinsic factors. Next, we asked where cognition is located. Here, internalism claimed that cognitive states are realized or instantiated wholly within the individual. Finally, we asked how cognition operates. The internalist response was that cognitive states are causally potent, manipulable internal information-bearing structures. The uptake is this: internalism makes three interrelated suppositions about the methodology of psychological practice, the location of cognitive processes, and the format and character of how cognition operates. Call these internalism's "tripartite assumptions".

At this point, it will be beneficial to offer three clarifications about the foregoing account. First, the account is not intended to be an exhaustive analysis of internalism.[3] Its aim is not to provide necessary and sufficient conditions for internalism. Rather, it is meant to capture the flavour of internalist accounts. Thus, the three features identified are best thought of as hallmarks of internalism. The idea is that highlighting the paradigmatic elements of internalism can draw attention to how such elements find their way into music cognition research. Second, each assumption should be read as forming a constraint on psychological investigation; they are limitations on the possible space of available explanations. I will say more about this in section 1.3.1. Third, in framing my analysis in terms of "information-bearing structures", I have sought to capture what's common to several approaches of internalism, while staying as neutral about the format and nature of the information-bearing vehicles as possible.[4]

---

[3] For further discussion of internalism/individualism, see Wilson (1995, 2004).

[4] The reason for this is that internalism has not always been clear about what position it should adopt with respect to computationalism, representationalism, and folk psychology (see Wilson, 1995, ch.2). Since these terms are technical terms of art in the philosophy of mind and cognitive science, I will try to avoid discussion of them where it is not directly relevant.

1.2 Internalism and Music Cognition

As with any subarea of psychology, music cognition research has many strands, each with its own rich history.[5] In this section, I argue that music cognition research has been committed, either implicitly or explicitly, to the tripartite assumptions of internalism. I attempt to establish this claim by looking at three exemplary projects of music cognition research.

However, one might reasonably ask why we should prefer this "exemplar" approach to a more systematic survey. There are at least two reasons. First, psychological practice is rarely explicit about its philosophical commitments (Chomsky, 1995). A fuller exposition will offer the depth needed to pull out the internalist assumptions where an otherwise more superficial survey might gloss over them. Second, each of the projects represents a larger stream or program of research that has been pursued. Selecting representative projects allows a larger swath of research to be covered. I will go into further detail about this point at the end of the section.

The projects reviewed here are: (i) Lerdahl and Jackendoff's generative theory of tonal music, (ii) Carol Krumhansl's perceptual theory of musical structure, and (iii) Michael Dawson's connectionist chord classifying networks. The projects here are by and large concerned with Western tonal-harmonic music. This fact should not be seen as a limitation on the current analysis, however. Rather, it is an attempt to capture what has been at the heart of most music cognition theory and research; it just so happens that traditional

[5] For a more comprehensive overview of music cognition/perception research see Radocy & Boyle (2012) or Deutsch (2012).

research has focused on a certain kind of tertian music (see Radocy & Boyle, 2012).

1.2.1 A Generative Theory of Tonal Music

The first project to consider is Fred Lerdahl and Ray Jackendoff's generative theory of tonal music (1983, 2006). Lerdahl and Jackendoff attempted to define the organizing principles behind musical cognition: "We take the goal of a theory of music to be a formal description of the musical intuitions of a listener who is experienced in a musical idiom" (1983, p. 1). Now, because Lerdahl and Jackendoff's project was inspired by Leonard Bernstein's (1976) lectures on applying Chomskyan linguistics to music, drawing a parallel with Chomskyan linguistics will help us get a better sense of the project's structure.[6]

According to the Chomskyan view, language acquisition in children faces a "poverty of stimulus" problem: The information available when first learning a language is in principle consistent with many different natural grammars. To explain how children deal with this problem, the Chomskyan approach posits that children possess tacit or innate knowledge of syntactic rules, what Chomsky calls the "universal grammar". As Chomsky says, "[t]he child approaches the data with the presumption that they are drawn from a language of certain antecedent well-defined type" (Chomsky, 1965, p.27). This universal grammar or language faculty provides the additional information required to

---

[6] I should mention that in addition to being a theory of music cognition, Lerdahl and Jackendoff's account is also a theory of music. This is because part of the task of characterizing how people process music also involves laying out what structures there are to be processed. But because my interest lies more with music cognition then with theories of music, I will largely ignore this additional dimension of the project. This point also applies to the Chomskyan analog.

learn particular languages. Instead of learning a whole grammar, children face the much more tractable problem of simply refining or adjusting sets of innate grammatical rules.

On the Chomskyan approach, the major task of linguistics is to identify the components of the universal grammar. For example, Chomsky (1965) posits two types of rules for the universal grammar: "substantive universals", which represent the fixed set of syntactic categories, and "formal universals", which represent the rules for the abstract properties every natural language possesses in order to be grammatical. The task of the linguist is to describe the syntactic rules and how they apply in order to explain linguistic competency.

It should be mentioned that Chomsky's view of universal grammar has (notoriously) changed over time. Although I have presented Chomsky's earlier "rules and representations" view (Chomsky, 1959, 1965), there was shift in Chomsky's thinking to a "parameters and values" view (see Chomsky, 1995). The choice to present Chomsky's earlier view is due to its being theoretically closer to Lerdahl and Jackendoff's project. There are two notable ways in which the Chomskyan approach influenced Lerdahl and Jackendoff's project.

First, Lerdahl and Jackendoff also worry that musical structures are too complex and sophisticated to be explained by environmental acoustic features alone, similar to one of the chief motivations for the Chomskyan paradigm. They write:

> [A] listener without sufficient exposure to an idiom will not be able to organize in any rich way the sounds he perceives...once he [the listener] becomes familiar with the idiom, the kind of organization that he attributes to a given

> piece will not be arbitrary but will be highly constrained in specific ways. (Lerdahl & Jackendoff, 1983, p.3)

As a result of such poverty of stimulus concerns Lerdahl and Jackendoff also posit the presence of musical grammars in order to account for music cognition.

Second, like the Chomskyan approach, Lerdahl and Jackendoff take the goal of music cognition to be the formal description of musical grammars and how they are used to create or generate the complex musical structures. “The early work [of Chomskyan linguistics]...took as its goal the description of ‘all and only’ the sentences of a language, and many were led to think of a generative grammar as an algorithm to manufacture grammatical sentences. Under this interpretation, a musical theory should be an algorithm that composes pieces of music” (Lerdahl & Jackendoff, 1983, p.6).

We now have a sense of the Chomskyan view of language and how it influenced Lerdahl and Jackendoff’s project. In what follows, I consider how Lerdahl and Jackendoff cash out their Chomsky-inspired generative approach.

Lerdahl and Jackendoff’s musical grammar consists of two types of rules: well formedness and preference rules. Whereas the well-formedness rules operate analogously to substantive universals of Chomskyan linguistics, the preferences rules operate differently from the formal universals, as Lerdahl & Jackendoff themselves note (1983, p.5). Lerdahl and Jackendoff describe four different, but related, component hierarchies. Each of the hierarchies is associated with a set of well-formedness rules. First, there is a “grouping structure”, which segments a piece into motives, phrases, and sections; second, there is a “metrical structure”, which describes how sound events are associated with strong or weak beats; third, there is “time-span reduction”, which assigns structural importance to pitches relative to their position within the grouping and metrical structures; and fourth, there is a “prolongational reduction”, which assigns pitches an order of harmonic and melodic tension and relaxation, continuity and progression (Lerdahl and Jackendoff, 1983, p.8-9).

The rules associated with each of these hierarchies impose constraints on how musical structures can be generated; these are analogous to the constraining effects of the substantive and formal grammars of Chomskyan linguistics. For example, one well-formedness rule for the grouping hierarchy says that musical structures are organized so as to avoid groupings that contain a single event. This rule delimits some of the possible ways that a musical idiom can be psychologically structured (Lerdahl & Jackendoff, 1983, p.43).

However, the well-formedness rules alone are insufficient to parse musical stimuli. Unlike sentences, even when specified under some structural description, musical structures are neither correct nor incorrect. Rather, at best they can be said to be consonant or dissonant. Even if musical stimuli are organized by the well-formedness rules, this doesn't fully explain how individuals can have different parsing for the same musical piece. This leads Lerdahl and Jackendoff to supplement the well-formedness with preference rules. The preference rules specify the probable structures that listeners organize music stimuli with; they define an assignment of order within a listener's experience. Well-formedness rules answer "what-is-possible?" questions; preference rules answer "what-corresponds-to-expectancies?" questions (Boyle & Radocy, 2012, p.111).

To illustrate, consider the fact that musical stimuli are naturally grouped into perceptible chunks; the visual analog is the automatic partitioning of the visual field into objects, parts of objects, parts of parts, etc. How is this musical grouping accomplished? Lerdahl and Jackendoff's answer is that listeners use the grouping well-formedness rules to define the formal structures of the grouping patterns and then use the grouping preference rules to establish which of the formally possible structures are assigned to a piece (1983, p. 36-7). Lerdahl and Jackendoff lay out five well-formedness rules and seven preference rules for grouping, an example of which is provided in Figure 1.1.

Grouping Well-Formedness Rules (GWFR)	Grouping Preference Rules (GPR)
GWFR 1: Any contiguous sequence of pitch-events, drum beats, or the like can constitute a group, and only contiguous sequences can constitute a group (p.37).	**GPR 1**: Avoid analyses with very smaller groups (p.45).
GWFR 2: A piece constitutes a group (p.38).	**GPR 2 (Proximity)**: Consider a sequence of four notes n_1, n2, n3, n4. All else being equal the transitions n_2 - n_3 may be heard as group boundary if: (a) the interval of time from the end of n_2 to the beginning of n_3 is greater than that from the end of n_1 to the beginning of n_2 and that from the end of n_3 to the beginning of n_4 or (b) the interval of time between the attack points of n_1 and n_2 and that between the attack point of n_3 and n_4 (p.45).

Figure 1.1 Two well-formedness and preference rules for the grouping structure.

According to Lerdahl and Jackendoff, the well-formedness rules offer two grouping configurations of Mozart's Sonata K. 331 (Figure 1.2). Call them "grouping A" and "grouping B". While listening to a performance a listener can emphasize either grouping. This is because, as far as structural descriptions are concerned, each grouping is well formed. From this it follows that any difference between the two groupings from the listener's perspective must be due to a difference in preference rules. "These subtle variations in articulations are typical of the strategies used by performers to influence perceived grouping...in large part these strategies [groupings] are learned and used unconsciously" (Lerdahl & Jackendoff, 1983, p.64).

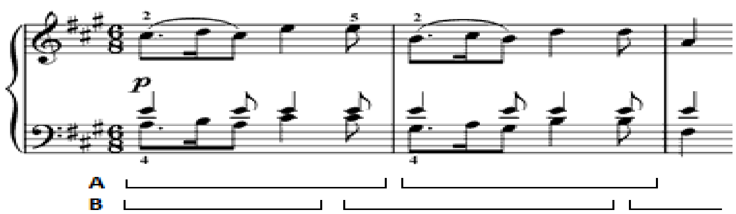

Figure 1.2 Mozart's Sonata K. 331 with two GWFR

To summarize, Lerdahl and Jackendoff's story goes something like this: First, listeners structure musical events using formal well-formedness principles; second, because of past experiences, preference rules are deployed to decide between the various configurations structured by the well-formedness rules. In short, musical cognition involves the organization and projection of structures onto sound events using formal principles.[7] Though this brief exposition probably hasn't done justice to the depth or scope of Lerdahl and Jackendoff's theory, because the aim is only to evaluate whether the project is committed to internalist assumptions, it should suffice present purposes.

[7] It's worth mentioning that Lerdahl and Jackendoff's account is largely theory driven. The inferences move from theory to data. Some have criticised this account on grounds of its psychological implausibility (Child, 1984).

Consider the theory's relationship to each of the tripartite assumptions. First, note Lerdahl and Jackendoff's motivation for their account:

> Rarely do two people hear a given piece in precisely the same way or with the same degree of richness. Nonetheless, there is normally considerable agreement on what are the most natural ways to hear a piece. A theory of a musical...should characterize situations in which there are alternative interpretations. (1983, p.3)

If two individuals are in the same environment, then any difference between their musical perceptions must be due to some difference in the underlying mental constructs. A psychological theory of music needs to make sense of "interpersonal agreement" and "alternative interpretations" cases. Therefore, the real interest of a musical cognitive theory lies in isolating the underlying mental structures and processes responsible for generating behaviour.

This approach is strikingly similar to the methodological assumption of internalism. Holding that environmental variables need to be kept constant implies that the musical idioms are relevant psychologically only insofar as they allow inferences to be made about which principles of organization are influencing an organism's internal processes.

Second, Lerdahl and Jackendoff hold that "insofar as one wishes to ascribe some sort of 'reality' to these structures, one must ultimately treat them as mental products imposed on or inferred from the physical signal" (Lerdahl & Jackendoff, 1983, p.2). Since physical signals are external to an individual, the musical structures must be "mental products" internal to the individual. The lack of direct correlates in the score or sound waves to musical structures suggests to Lerdahl and Jackendoff that the musical structures must be instantiated within the individual. This commitment is more fully expressed in a later article. They write: "The musical capacity constitutes the resources in the human mind/brain that make it possible for a human to acquire the ability to understand music in any of the musical idioms of the world, given the appropriate input" (Lerdahl & Jackendoff, 2006, p. 35). Talk of the "mind/brain" highlights quite clearly the assumption that the musical structures supervene wholly on internal, physical states of the individual.

Finally, given that the well-formedness rules and preference rules are the means to uncovering the mental entities, and that the mental products simply are the musical structures, it follows that Lerdahl and Jackendoff are committed to claiming that musical cognition involves the manipulation of information bearing structures or vehicles. Recall the poverty of stimulus concerns Lerdahl and Jackendoff adopted from the Chomskyan paradigm. They assumed that musical structures could only be generated when musical grammars were posited. This was because the musical grammars contained the information required to overcome the impoverished input. However, if the musical grammars combined with the incoming acoustic stimuli generate musical structures in order to account explain listeners' musical understanding, then the psychological musical structures that are internal to the individuals are information-bearing structures in the required sense of the operational assumption.

1.2.2 A Perceptual Theory of Tonal Structure

The second project to consider is Carol Krumhansl's perceptual theory of tonal music (1979, 1982, 1983, 1990). Krumhansl investigated pitch patterns processing, particularly in tonal contexts. The motivating question was how pitch patterns could be perceived as musical given that they are largely unstructured. Krumhansl posited that musical structures (e.g., pitch relationships) are encoded, remembered, and then compared to incoming stimuli.

Krumhansl's theory is supported by an experimental paradigm called the "tone probe method" (Krumhansl & Shepard, 1979). According to this method, a musical context is established, for instance by playing a partial scale or a chord. After this, a probe note is played. A subject then rates the probe according to how well it fits within the musical context, e.g., by rating how well it completes a partial scale. The method requires that listeners make judgments about the relatedness of musical elements, usually on a numerical scale. The design of task has the potential to reveal how musical elements depend on context key and how elements conform to a listener's experience.

To illustrate, consider two tone probe experiments. First, in Krumhansl and Shepard (1979), subjects listened to a C major scale or triad, followed by a pair of tones; the tone pairs were randomly selected from the C major scale. When subjects rated the tonal similarity of these tones it was found that tones from the C triad (C-E-G) were heard as more similar than tones from the C major scale; particularly this was the case for nondiatonic tones (tones other than scale tones). Second, the same judgment task was used, but the stimuli were changed. Subjects heard a standard tone (the tonic) and then a pattern of eight interpolated (inserted) tones. These tones were either tonal or atonal. This pattern was then followed by sounding a final comparison tone. The task was to judge the similarity of the standard tone to the comparison tone by rating how closely they matched in spite of interfering effects of the interpolated tones. It was found that it was easier for subjects to remember the diatonic standard tones. The explanation offered was that it seemed that the interference was less significant for the diatonic interpolated tones than for the nondiatonic tones.[8]

Krumhansl interpreted such studies as revealing that judgments about similarity, whether with or without interference, tend to group around the tonic. This led her to hypothesize that within any given musical context individual pitches are perceived inside a hierarchy of tonal organization. A tonal hierarchy describes the relationship of stability between musical pitches for a given pitch class. Because musical pitches can be said to vary in terms of their relatedness to (i.e., consonance or dissonance from) an established central tone or tonic, they can be hierarchically organized as more or less stable.

[8] The subject populations of Krumhansl's studies varied. Some used individuals who had some musical exposure, usually in the form of experience playing an instrument; others used individuals had little formal music theory training.

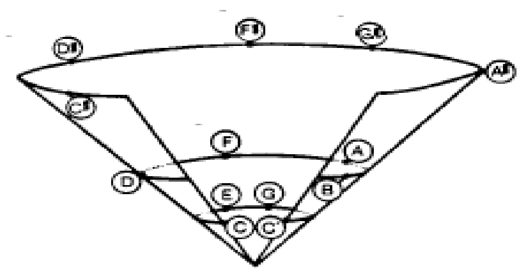

Figure 1.3. Krumhansl’s tonal hierarchy

To see how Krumhansl arrives at this conclusion, consider the first experiment again. For any particular trial, the key of C major establishes the key context. This means that the C is the central tone or tonic. In music theory, the closest stable tones to the tonic in C major are C-E-G. Given that subjects tended to rate those tones as the most similar to the C, this suggests that the subjects’ judgments are informed in some way by knowledge of the stability relations between tones. “Tonal organization resides in the functioning of the set of musical pitches around a single pitch, the tonic, may reflect a general principle of perceptual-cognitive organization” (Krumhansl, 1983, p.37). It is because the tonal hierarchy is used by listeners in a given context to structure perception that listeners make the judgments of tonal music they do. After extensive use of the tone probe method, Krumhansl placed the tonic triad at the top of the tonal hierarchy and diatonic and nondiatonic tones further down -- the hierarchy can be represented as a three dimensional cone (Figure 1.3).[9] With a sketch of Krumhansl’s theory in hand, consider how it relates to the tripartite assumptions of internalism.

---

[9] It bears mentioning that Krumhansl’s theory is largely an inductive one. The inferences about music cognition move only from data to theory.

First, consider Krumhansl's explanation of her methodology: "It is essential to control irrelevant properties of the stimulus, and to change or manipulate only those of interest. Potentially separable properties must not be experimentally covaried or confounded if the pattern of response is to be interpreted unambiguously" (1990, p.8). There are two points to note here. First, the irrelevant properties are external to listeners; second, the "interpretation" requires external properties to be held constant. The only reason why these two features should be in Krumhansl's methodology is if she accepts that environmental factors are psychologically relevant only insofar as they influence an organism's current, internal physical state. "[U]nder what are seemingly identical external circumstances, different individuals will respond differently, and the same individual will responds differently on different occasion" (Krumhansl, 1990, p.8). However, this is the methodological assumption of internalism, as it a version of the principle of autonomy.

Second, note how Krumhansl conceives of the role of the tonal hierarchy in (tonal) music perception: "[L]isteners, at least those with a moderate level of experience with tonal music, have apparently internalized this aspect of musical structures and use this knowledge to encode and remember pitches" (Krumhansl, 1983, p.43). The important point to note is that the tonal hierarchy is not a musical property but, rather, a psychologically imposed organization of musical elements (Dawson, 2013, p.275). Because the sound patterns are too "improvised" and "unorganized", the hierarchy is internal to the listener. However, notice that internalization entails that musical structures are realized by wholly in-the-head physical states. This is why Krumhansl writes, "[t]he aim is to describe the human capacity of internalizing the structure sound materials of music by characterizing the nature of internal processes and representations" (Krumhansl, 1990, p.6). This is the second metaphysical assumption of internalism.

Finally, recall that the tonal hierarchy describes relationships of tonal stability. Relevant here is that Krumhansl views the tonal hierarchy as a "mental representation" (1983, p.2). It is because the tonal hierarchy is a "scheme" or "template" that it can be compared with incoming stimuli that judgments about tonal music can be made. But if this is true, Krumhansl's view is committed to the third assumption of internalism: It posits the tonal hierarchy as an internal, information-bearing structure.

1.2.3 Connectionism and Music Cognition

So far I have concentrated on projects that attempt to explain music perception by appealing to highly structured information bearing vehicles: Krumhansl's geometrical tonal hierarchy and Lerdahl and Jackendoff's musical structures. However, not all music cognition research has employed such discrete, symbol-like structures in order to account for music cognition. Connectionist researchers have opted for more "distributed" vehicles. It will be worthwhile to explore an example of this kind of research.

Connectionist or artificial neural networks are computer simulations of "brain-like" systems. They are made up of two elements: processing units and connection weights. As the processing units become active, either by turning on or off, they transmit signals to other units through the weighted connections. When the signals reach and activate the output layer of units the network has generated a response. The set of all the signals between the processing units represents the "knowledge" of the network. Because networks usually have small, randomly assigned connection weights at the start of a simulation, they are unlikely to initially generate the correct response. As the networks correct their outputted responses across trials, by employing "learning rules", they are "taught" the correct response. When the output response finally matches the desired activity, the network has solved its problem. As an example of such a network, consider Dawson et al.'s (2008) study of a multilayer perception trained to classify different types of musical chords.

In Dawson et al. (2008) a multilayered perception was trained to classify four chord types: major, minor, dominant seventh, and diminished seventh. The network had four output units, three hidden units, and 12 input units. Each of the output units represented one of the four types of musical chords that the network needed to classify, while each of the input units represented a particular musical note in a pitch class (Figure 1.4). The input to the network was a set of 48 different chords, constructed by building each chord type on the root note of the pitch class; for example, the C major chord used the C note as its root and then activated the three input units that defined the component notes of the chord: C, E, and G.

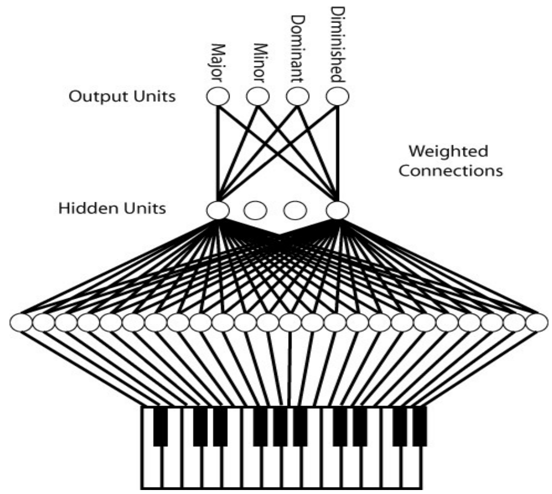

Figure 1.4 Dawson's Chord Classifying Network

Dawson et al. found that this type of network was able to successfully classify all of the chords in the training set. Importantly, it was found that the network was able to classify the chord types because the connection weights "assigned" to the input units three "note names" that corresponded to one of four equivalence classes. These equivalence classes used different musical intervals, for example the circle of major thirds or major seconds, to classify chords by selectively responding to incoming stimuli (Dawson & Yaremchuk, 2008, p.28).

To illustrate, consider how the network classified a diminished seventh tetrachord -- a four note chord, e.g., C, D#, F#, A. When a diminished seventh tetrachord was presented to the input units, the connection weights that coded for the equivalence class activated either Hidden Units 1 or 3 (Dawson & Yaremchuk, 2008, p.26). What this means is that the classification of the diminished seventh tetrachord was the product of the hidden units sensitivity to a variety of feature values, not just particular feature types, features that were encoded in the connection weights. Dawson et al.'s interpretation maintained that the network abstracts musical structures in the form of musical intervals in order to classify chord structures. In this way, the "representations" were distributed across the connection weights; they were not encoded or localizable to any specific hidden unit. The explanatory import of this research is that it offers a possible model to study how people might classify chord structures (Dawson & Yaremchuk, 2008, p.28). The question is how this type of research relates to the internalist assumptions.

First, notice how well Dawson et al.'s study comports with the methodological assumption of internalism. The input representations are taken to be impoverished: Each chord is made up of only three to four notes corresponding to the "notes" of the input units. Because the network is taken to use elaborate internal structures to classify the chord structures, the inputs are taken to be of secondary importance. "[T]he network classified chord structure first by representing individual notes in terms of circles of major thirds and major seconds, and then by combining these representations to position chords in a three-dimensional hidden unit space" (Dawson & Yaremchuk, 2008, p. 28). Recall that methodological solipsism holds that psychological states should be construed without reference to anything beyond the boundary of the individual. If the inputs are the outward boundary of the "individual", then the real locus of explanatory relevance must be the hidden units and connection weights: a clear statement of the methodological assumption.

Second, consider that Dawson et al. focus heavily on interpreting the internal structure of the network: “One of connectionism’s potential contributions to the psychology of music is its ability...to suggest new approaches to represent musical patterns. In order for this potential to be realized, it must be possible to analyze the internal structure of a network after it has been trained” (Dawson & Yaremchuk, 2008, p. 28). How does network interpretation work? Dawson (2009) claims that “[t]he interpretation of the network is the process of assigning representational content to internal network components and to network states. Furthermore, one appeals to these representational contents in an attempt to account for the behaviour of the network” (p.185). However, if, as we saw, the network’s architecture consists only of processing units and connection weights, then focusing on analysing the network’s internal structures for representations only makes sense if the internal components being ‘assigned’ the representational content (i.e., bearing information) constitute the representations’ realization base. For if they didn’t, then such talk of the representational states as causally productive wouldn’t be coherent. Although implicit, Dawson et al.’s focus on network interpretation reveals a larger commitment to the internalist supervenience thesis, i.e., the metaphysical assumption.

Finally, Dawson et al. point out that the representations used by the network to solve the classification problem are distributed. “To our knowledge, the only previous occurrence of this kind of representation was the distributed coding scheme” (Dawson, 2008, p. 28). No single hidden unit is responsible for representing any one-chord type. Though significant that the information bearing structures are different from those of other music cognition studies, it is still the case that the causally potent vehicles are internal. The distributed representations are the information-bearing states of cognition. Dawson et al.'s study is committed to the operational assumption of internalism. In spite of some apparent differences, connectionist research remains quite faithful to the tripartite assumptions.

1.2.4 Underdetermination, Internalism, and Music Cognition

I have argued that three programmatic projects of music psychology are committed to the tripartite assumptions of internalism. As was mentioned earlier, though, one concern might be that review of three projects is insufficient to show that music cognition more generally has a deep commitment to internalism. Perhaps all that has been shown is that a small section of outlying research has made particular bedfellows. To show the stronger claim that music psychology more generally has been committed to internalism, the tripartite assumptions need to be shown in a larger sample of research.

Although I have some sympathy with this concern, by noting that each project represents a different research strand in music psychology, it can be allayed. To see how this might be the case, attention can be drawn to the role of underdetermination concerns in structuring music psychology research.

Theories in cognitive science are often taken to face problems of underdetermination (see Dawson, 1998, 2013). An underdetermination problem occurs when information does not uniquely determine a correct solution. One example from the study of visual perception is that the information presented to the eyes (the proximal stimulus) is consistent with an infinite number of different scene interpretations (models of the world). And yet, only one interpretation is correct for almost any given context. To solve this problem, further knowledge about the world is required. Perception researchers offer various hypotheses about how the additional knowledge is obtained by the visual system. Some, for example, maintain that the additional knowledge is present innately, wired into the visual system itself (Pylyshyn, 2009), others that it is provided by the world in the form of natural constraints (Marr, 1982).

Concerns of this stripe have also motivated psychological theories in music psychology. One way to see the history of music psychology is as attempts to answer underdetermination concerns. Three particular answers stand out, each of which corresponds to one of the three projects.

First, for some, underdetermination concerns require invoking the presence of "internal frameworks". According to this tradition, abilities to perceive, remember and act on musical information stem from the formation and use of schemes or templates (West, Howell, Cross, 1985; Cuddy, 1991; Bulter, 1989; Dowling & Harwood, 1986). As we saw, Krumhansl's perceptual tonal theory posited such internal frameworks in the form of the tonal hierarchy. The incoming musical stimuli were compared to the tonal hierarchy. This allowed subjects to make tonal judgments.

Second, similar to how generative linguistics took children to face a poverty of the stimulus problem when first learning language, others have thought to overcome the impoverished nature of musical stimuli by appealing to innate, organizing principles (Temperlery, 2001; Hammanaka, Hirata, & Tojo, 2006). It was in this vein that Lerdahl and Jackendoff posited the musical grammars. The well-formedness and preference rules of the musical grammars constrained and described musical structures, supplying the additional information to overcome the initial poverty of stimulus.

Finally, connectionism thinking has influenced some researchers. To address underdetermination concerns, some researchers have appealed to the distributed representations found in connectionist networks (Bharucha, 1999; Fiske, 2004; Griffith & Todd, 1999; Todd & Loy, 1991). Through a process of feature extraction, elaborate internal representations are constructed, supplying the requisite additional information. As was quite clearly shown, Dawson et al.'s study opted for just this approach when it used a connectionist network to explore chord classification.

As is hopefully clear from this and the previous discussion, each project is motivated by underdetermination considerations, the result of which is the positing of rich, internal structures. Such considerations should lead us to view that each project has quite a bit in common with the traditional approaches to research. Each project captures certain features distinctive of the three research traditions.

This discussion should also help highlight in part why music psychology has had an allegiance to internalism. For when the world is seen as a swirling, chaotic mess of stimuli, as it is sometimes presented in underdetermination concerns, it is easy to look at internal cognition as imposing order, whether that is through schemes, innate grammars, or distributed representations. There is a rather clear connection, then, between underdetermination concerns on the one hand and internalism on the other.

1.3 Internalism, Explanation, and Music Cognition

In the previous section, I showed how three exemplar music cognition projects, and by extension music psychology more generally, have been committed to the tripartite assumptions of internalism. Now I want to consider the implications of the internalist commitments for music cognition research. I will argue that, on occasion, an internalist music psychology has less explanatory power or depth than music psychology that takes a wide approach. On route to this conclusion, I want to first relate psychological explanations to internalism.

The dominant form of explanation in psychology is "mechanistic explanation" (Bechtel, 2007; Wright & Bechtel, 2007; Cummins, 1983, 2000). According to this view, psychological explanations attempt to identify the mechanisms that are casually responsible for producing psychological phenomena. Mechanisms are sets of components that are organized such that they produce an activity or function (Wright & Bechtel, 2007, p.119). Psychological researchers attempt to identify the salient components or properties that coordinate to form the composite system that produce psychological phenomena. In practice, mechanistic explanations often involve representing or modeling the mechanisms in order to make inferences about how the mechanism supports the phenomenon of interest.

Do psychological explanations have to take a mechanistic form? Isn't it possible instead that they involve appeal to laws? Perhaps a mature approach to psychological explanation requires subsumption under laws, where the explanandum has to be expected as the conclusion of a sound argument (see, e.g., Hempel, 1965).

Cummins (2000) offers two related reasons to think that this isn't the case. First, laws in psychology are essentially only specifications of effects. For example, the primacy and recency effect describes the fact that subjects tend remember items at the beginning and ends of lists better than those in middle. Yet a law describing data is not the same thing as explaining it. Just because data confirms some effect described by a law does not explain the cause of why people tend to remember the first and last items of lists better. Laws only describe what cognition effects; they do not explain how it works (Cummins, 2000, p.120).

Second, for the above reasons, laws do not target the real explanandum of psychology. The primary explanantia of psychology are capacities or functions, e.g., the capacity to perceive depth, to speak a language, to hear music. It is capacities that cause the effects (i.e., behaviours) described by the laws. They are the explanantia of the behaviour described by laws. Identifying and describing the mechanisms responsible for behaviour is real task of psychology. “Psychology should seek to discover and specify the effects characteristic of the systems that constitute their proprietary domains, and to explain those effects in terms of the structure of those systems” (Cummins, 2000, p.122). There seems to be good reason to think that psychological explanations should be construed as the search for mechanisms.

Interestingly, nothing in principle bars mechanistic explanations from identifying composite mechanisms that include components internal and external to the organism or individual. Call individual-bound mechanisms “narrow” and world-traversing mechanisms “wide”. Several authors have, in fact, acknowledged the possibility of wide mechanisms; see Menary (2007, ch.1) on “wide capacities”. What I want to do here is draw attention to how wide and narrow mechanisms figure into psychological explanations. It seems that the “width” of psychological explanations, in part, depend on the location of the causally operative components that form the composite mechanisms. If the explanatory focus is put on intrinsic components, the explanations will be narrow, since the identified mechanisms are individual bound. But if attention is drawn outward to external components, explanations can identify mechanisms that spread across external and internal parts.

To illustrate, consider that one could explain muscle contraction by describing a mechanism that includes only intrinsic properties of the organism; for example, the coordinated firing of a group of neurons. In this case, the explanation would be narrow, because the supporting mechanism would be formed by elements contained wholly within the organism. However, if one explained, for example, an organism withdrawing into its shell by invoking both internal and external components, say, muscle contraction and the presence of predators, then the explanation would be wide, because the explanation would identify a mechanism that spreads across organism and environment. The uptake is that the width of mechanistic explanation depends on the causal properties (either wide or intrinsic) that it identifies as crucial for the relevant mechanisms supporting the activity of interest.

Does this mean that wide and narrow explanations cannot address the same phenomenon? Not necessarily. Depending on contextual features, the same phenomenon can be given either a wide or narrow gloss. For example, if investigated by a physiologist, withdrawal behaviour might be described solely in terms of electrochemical mechanisms, e.g., properties of neurons firing. However, if investigated by an evolutionary biologist, one interested in the larger effects of the ecological niche, that same withdrawal behaviour might be explained by appeal to an integrated mechanism that includes environmental components, e.g., the presence of predators, and internal properties, e.g., neuronal firings. The salience, and therefore identification, of the coordinated subcomponents of mechanisms is influenced by larger explanatory goals. I will return to this point in section 1.3.1. For now the question is how wide and narrow explanations relate to internalist psychology.

1.3.1 Internalism and Narrow Explanations

According to internalism, psychological explanations are and should be exclusively narrow (Menary, 2007; Wilson, 1994a, 1995). Consider Fodor's

(1987) "argument from causal properties". Wilson (1995) gives the argument the following form:

(i) In the cognitive sciences, both mental causes of an individual's behaviour and that behaviour itself must be individuated in terms of causal powers of that individual.

(ii) The causal powers of anything are determined or fixed by that thing's intrinsic physical properties

(iii) Therefore, the cognitive sciences, particularly psychology, should concern themselves only with states and processes that themselves are determined by the intrinsic physical properties of the individual.[10]

Suppose that the argument were sound. What follows? First, it would seem that if premise (i) were true, then explanatory practice would have to be concerned only with internal components, wide or external components would only be of secondary importance. Second, if premise (ii) were true, the mechanisms would be identified with local neuronal structures. This would mean that the salient subcomponents would have to be intrinsic to the organism, and so the explanations narrow. If the argument from causal properties were sound, it would seem to entail narrow explanations. It seems fair to say that narrow explanations can be derived from the internalist construal of cognition (see Wilson, 1995, ch.2). It will be worthwhile to dwell a bit longer on this question of how narrow explanations are constrained by internalism.

[10] This argument has been extracted from a longer argument presented in Wilson (1995, p. 32)

Consider again the tripartite assumptions. The tripartite assumption can be said to constrain music psychology explanations by delimiting the kinds of components or properties that can figure into the supporting mechanisms. The methodological, metaphysical, and operational assumptions restrict the theoretical horizons of what resources count as cognitive for psychologists; or, in the present discussion, music psychologists. Consider an illustrative story.

While Willie Sutton was in prison for robbing banks, a reforming priest asked him why he robbed banks. He replied: “Well, that’s where the money is”. Clearly, there is a failure between the Sutton and the priest to connect on what the relevant answer to the priest’s question is. Garfinkel (1981) represents the situation like this:

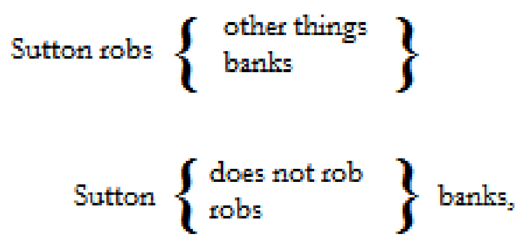

Figure 1.5 Constraints on answers

Why is there a failure between Sutton and the priest? The obvious but informative answer is that the background purposes and values constrain the answers each see appropriate: the priest by his interest in reform, Sutton by his desire to rob banks (see Garkinfel, 1981, p.22). Because of the background assumptions, the space of alternatives is restricted. Sutton takes the space of possible answers to include the banks he robbed and other banks, while the priest takes the contrast space to be between robbing banks and not robbing banks at all. An analogous situation is at play for the tripartite assumptions and narrow music cognition explanations.

The tripartite assumptions delimit the space of alternative explanations for music psychologists by restricting the kinds of properties that can properly figure into the mechanisms identified by explanations. The explanations of musical phenomena provided by people like Krumhansl or Lerdahl and Jackendoff take on a narrow form because the viable mechanisms that can be identified can only include components contained within the individual. This is in part why Krumhansl posits the tonal hierarchy, Dawson appeals to the network's distributed representations, and Lerdahl and Jackendoff look to musical structures generated by the music grammars.

Like the option between robbing the bank he robbed and other banks, music cognition explanations are constrained in the type of mechanisms they can appeal to. Conversely, this also means that wide mechanisms, and therefore wide explanations, do not figure into the contrast space; they do not figure into musical psychologists explanatory horizons as plausible alternatives. To borrow from Garfinkel, the situation can be represented as such:

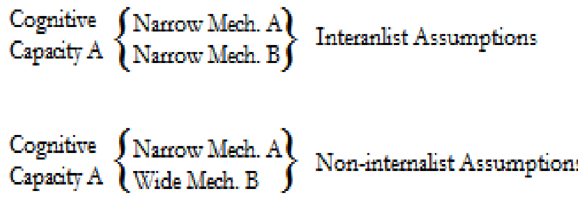

Figure 1.6 Constraints on psychological explanation

Given this state of affairs, I want spend the remainder of the chapter making space for wide explanations.

1.3.2 Explanatory Power and Music Cognition

If music cognition research has made notable use of narrow explanations, the question to answer is whether this is desirable. This pushes the discussion into the realm of explanatory "power" or "depth". To know what it means to say that one explanation is better than another is to look for the properties that identify explanations as better or worse.

To begin, notice that explanations need to share the same explanatory task in order to be compared. This makes sense given that if they didn't, they would be explaining different objects, and thus really not be standing in comparison.[11] So, assuming explanations have same explanatory task, to say that one explanation has more "power" than other is to say that it allows more appropriate counterfactual inferences (Wilson, 1994a; Ylikoski & Kuorikoski, 2010). Or less technically, an explanation has power insofar as it possesses some properties that improve its ability to enhance understanding. One

[11] There are other instances where comparisons are possible, but this is the most relevant condition for present purposes. For other conditions see Ylikoski & Kuorikoski (2010).

advantage of this formulation is that it offers a link between explanation and understanding. The more counterfactual inferences an explanation allows the more "what-if-things-had-been-different" questions it answers and therefore the more information it provides.

In the psychological context, this means that mechanistic explanations are more counterfactually robust if they identify mechanisms that hold across a greater variety of circumstances. For example, in the Krumhansl case, the tonal processing mechanism, which operates on the tonal hierarchy, is supposedly the most counterfactually robust mechanism, because subjects reliably produce tonal behaviours (i.e., make judgments) that are consistent with its presence.

This account of explanatory power can be distinguished from what are sometimes called "evidential virtues" (Ylikoski & Kuorikoski, 2010). While explanatory power looks to provide evaluations according to explanatory desiderata or virtues, evidential virtues look at the likelihood of explanations being true given some evidence. Explanatory virtues refer to how good an explanation is if it is true, evidential virtues refer to how much better an explanation is supported by evidence over a rival. In practice, both these types of virtues are evaluated together. In the present discussion, the interest is in the former not the latter. Consider two aspects of explanatory power.

First, explanations can be said to be more or less "theoretically appropriate" (Wilson, 1994a; 1995; Ylikoski & Kuorikoski, 2010). Explanations can describe phenomena at various levels. For example, a coffee mug can be simultaneously described microphysically as a mass of molecules and macrophysically as a mass of ceramic. The truth of one level of description does not exclude the truth of another. However, the truth of a description is not the same thing as the appropriateness of an explanation. Sometimes it is more appropriate to pitch characterizations at one level rather than another.[12]

For example, to explain a failed interview it seems more appropriate to appeal to the theoretical resources of the social and economic sciences rather than those of neurology, even when they both make the same predictions. The intuition is that it is desirable to have explanations that match the levels at which phenomena are most naturally characterized. This desideratum also holds true of explanations in music cognition. As we have seen, internalist music cognition favours narrow explanations. However, if it turned out that some musical cognitive phenomena were more appropriately characterized at levels amendable to wide explanations, there would be reason to explore non-internalist approaches to research.

To motivate this possibility, consider Balzano's (1986) investigation of the effects of structural pitch-time constraints on music identification. Balzano had participants listen to two kinds of pseudomelodies and then judge whether they were more or less "musical". They either heard pitch pairs that used identical rhythm (time structure) and different pitch constraints or time pairs that used identical pitch and different rhythm (time structure). When compared across various conditions, listeners were significantly more responsive, in terms of

[12] Either fortunately or unfortunately, this talk of "levels" will be a common theme throughout the thesis.

discriminating sounds as musical, to those conditions that included the "invariant" properties of the pitch-time constraints.

The question is what the most theoretically appropriate level at which to characterize this phenomenon is. The response from internalist music cognition would be to offer a narrow explanation via an internal, information-processing mechanism(s) (see Balzano, 1986, p.216). However, such a gloss would seemingly fail to reflect the coarseness of the behavioural responses; the fact that listeners respond directly to the presence or absence of the informational structures in the sound itself. It will fail to provide the optimal amount of information about the target phenomenon, because it would be less able to account for local variation. Explanations that are either too acute or obscuring are less theoretically desirable. If narrow explanations only reference internally constituted mechanisms, they obscure the important structuring nature of the invariant properties. Taking a wider approach to mechanism identification would preserve the sustaining role of the invariants. Balzano's research seems more appropriately characterized by a wide rather than narrow explanation.

Second, explanations can also be said to have more or less "causal depth" (Morton, 2002; Wilson, 1994a; Ylikoski & Kuorikoski, 2010). If the properties identified in an explanation are more counterfactually robust, the explanation has more causal depth. For instance, the temperature of the rotten vegetation in which the Mississippi alligator lays its eggs determines the sex of it offspring (Clark, 2008, p.7). Although there is a genetic component added by the alligator, explanations identifying the temperature are causally deeper than those that fail to, because they identify properties that are more resistant to change (i.e., changes in the surrounding vegetation correspond more robustly with the sex of the offspring).

To take a more cognitive example, consider passive dynamic walkers (PDW). PDWs are energy-efficient walking agents that couple their basic dynamic morphology to worldly walking apparatuses; walking is accomplished using only mechanical linkages and the properties of the environment (Clark, 2008, p.5). Any explanation that tries to explain the locomotion of the PDWs without referencing the worldly structures that generate and sustain their walking behaviour will be counterfactually less robust. This is because such explanations will be at a loss to account for how the walker fails to walk on surfaces not conducive to the natural dynamics of the agent. The worldly structures are deep causal determinants of the walking system.

Finally, consider an explicitly psychological example. It turns out that the people can retrieve the details associated with a previously experienced event (recollection), without feeling that the event was previously experienced (familiarity). And vice-versa. This has led to a debate about whether recognition and familiarity are supported by one or two memory processes. Two competing models have emerged: single and dual process models (see Wixted, 2007; Rugg & Curran, 2011). The debate centers on which model identifies the most causally robust properties. The dual-process theory claims that certain neurons in the hippocampus are responsible for familiarity memory, but not recognition. Single process theories respond by arguing that interpreted correctly, hippocampal neurons are completely responsible for both kinds of memory. The model that is right identifies the most counterfactually robust properties composing the mechanism(s) responsible for recognition and familiarity. It would have more causal depth than its rival. Explanations that identify more counterfactually robust mechanisms have more causal depth. What's more, if it can be shown that narrow explanations possess less causal depth than their wide counterparts, this would be to, again, motivate a non-internalist approach to music cognition.

So, consider Balzano's research again. To recall, the main result was that subjects identified those acoustics sequences as musical which had a greater degree of invariant pitch time constraints; what Balzano called generative and quantal properties. The invariant properties are deep causal determinants of the ability to identify music qua music. For this reason, any explanation that fails to account for them will be counterfactually less robust. Because of its focus on narrow mechanisms, ones that consist of intrinsic components, narrow explanations cannot make a place for the invariant properties in the explanation, except as inputs. Yet construing the invariant properties as inputs fails to acknowledge the operative, causal place they have in generating and sustaining subjects' behaviour. Without the invariant properties, there would be a notable drop in performance, i.e. music identification. Narrow explanations would not provide as casually deep an explanation, because they would fail to include the operative invariant pitch-time properties.

Wide explanations, on the other hand, would account for the operative components (the invariant pitch-time constraints), because the mechanisms they identify could include environmental structures. An explanation that identifies organism-environment crossing mechanisms would be more resistant to change, because it could include parts that are more stable across local variation. From this it follows that a wide explanation that identifies a wide mechanism, one which includes invariants, would offer a causally deeper explanation than its narrow counterpart.

One might reasonably ask, though: isn't it an open question whether the causal deep factors are in the sound or the internal representations? Could it not be the case that the perceivers are exploiting properties of internal representations rather than external structures? One reason to prefer thinking they are in the sound (besides the fact that Balzano himself claims that they are) is that the move to internal representations is unmotivated: Internal representations really only need to be invoked when stimuli are impoverished (see section 1.2.4). Since the invariant pitch-time constraints can be coherently described as external properties, ones that form rich, informational structures, the internal representations invoked by underdetermination concerns are unnecessary. Invariant properties supply the required information. We seem to have good reason to suppose that the deep causal factors are the invariant properties of sound itself, not internal representations. The uptake is this: wide explanations, on occasion, can offer explanatorily more powerful accounts of music cognition research.

So, here is the argument that has been developed:

(i) Music cognition research has been committed to an internalist paradigm of thinking.

(ii) Internalist thinking places constraints on the kinds of explanations music psychology offers.

(iii) Music psychology explanations informed by internalist thinking can have less explanatory power or depth than wider ones.

(iv) Therefore, on explanatory grounds, we have reason to explore the possibility of wide music psychology.[13]

1.3.3 Conclusion

Suppose for the moment that what's been said is on track, that on the two dimensions of explanatory power wide explanations are sometimes superior. What does this show? Surely it does not show that internalist music cognition should be abandoned wholesale. Hasn't the importance of these two dimensions of explanatory power been overemphasised? Not quite. Though it's right to say that two dimensions of explanatory power aren't enough reason to abandon internalist music cognition wholesale, they do provide enough reason to favour exploring an alternative approach. This modally weaker claim is more motivational than deconstructive.

[13] It is worth mentioning that this argument does not claim that it is because wide music psychology offers better *predictive success* than an internalist approach that it should be preferred; it is not an inference to best explanation. Rather, it argues that even if both approaches have the same predictive accuracy, there is still further *explanatory* reason to explore wide music psychology. In this way, the argument bypasses several recent objections to theories antithetical to internalist cognition based on underdetermination considerations (see, e.g., Baker, 2010; Sprevak, 2010).

As I have tried to show, narrow explanations can become strained. When this happens one should begin to explore alternative approaches, ones that adopt explanatorily richer wide explanations. Shifting to wide music psychology only requires giving up a global view about internalist music psychology, not the weaker thesis that internalist music cognition is informative and profitable. In this way, the argument that has been developed embraces a kind of pluralism (see, e.g., Wilson, 1994b, 2014). It is when music psychology's penchant for internalist-friendly narrow explanations obscures more powerful explanations that there is reason to explore an alternative approach. In the next chapter, I want to develop just such an alternative.

Chapter 2 - Wide Music Psychology

2.0 Chapter Overview

There are two goals of this second chapter: (i) to flesh out the details of wide music psychology and (ii) to show how music cognition is "extended". I begin by outlining the basic theoretical commitments of wide music psychology and some putative examples of its research. The aim is to provide a picture of wide music cognition and set the stage for exploring music cognition as an extended phenomenon. [14]

2.1 Externalism and Music Cognition

In Chapter 1 we saw that wide explanations can be explanatorily superior to narrow ones. This concession has an important consequence for theorizing: externalist music cognition needs to be able to support the use of wide explanations. Given this, it seems that there are three basic, structural assumptions wide music psychology needs to adopt -- each corresponding to one of the methodological, metaphysical and operational aspects of research. For lack of a better name, call these the "externalist assumptions". There are two consequences of adopting these assumptions: (i) they open up the conceptual space for wide explanations and (ii) they provide a picture of an analytical framework that wide music psychology might use.

[14] A version of this chapter has been submitted for publication. Kersten, L .(2014). Empirical Musicology Review. 9(2).

First, wide music psychology needs to be able to maintain that music cognition requires focusing less on what is occurring inside of the cognizing individual and more on external structures. Second, wide music psychology needs a metaphysics that is able to support the existence of wide mechanisms, given their central role in wide explanations. And third, wide music psychology needs to make room for external information-bearing states that can figure into wide mechanisms. Let's consider each in turn.

2.1.1 Looking Beyond the Individual

First, if wide explanations identify mechanisms that include components external to individuals, then the approach to studying cognition needs to be rethought. If the internal features are no longer the prime explanatory movers, then methodologically a new approach is needed that takes seriously the contributions of the environment. An approach that turns to understanding internal structures only after the cognitive role of external structures has been exhaustively explored. To support wide explanations, wide music psychology needs to be guided by the assumption that turning to internal processing can only happen once an investigation has been made into how organisms off-load and distribute cognitive processing via worldly structures.

Rowlands (1995) offers a good starting point. Rowlands, following Gibson (1966, 1986), describes two metatheoretical principles for the study of cognition: (i) the amount of internal information required for some cognitive system to accomplish a task is inversely proportional to the amount of information contained in the environment and (ii) we cannot understand how cognitive systems accomplish their tasks unless we know what information they obtain from the environment (Rowlands, 1995, p.5). These two principles signal the kind of shift needed for wide music psychology. They point to how psychological studies need to be first and foremost sensitive to the presence of external structures.

Wilson's description of "integrative synthesis" draws a similar moral. Talking about the physiological sciences, he writes: "here anatomically identified organs [are] located within some broader, functional system, with integration replacing location as the corresponding technique of investigation" (2004, p.32). Similar to Rowlands' principles, the idea is that the methodological gaze needs to shift outward beyond the organ to the larger physiological system. Of course, the larger system is still internal to the organism, but the methodological promise is the same: theoretical space needs to be made for external structures. In stressing function and integration over location, integrative synthesis reveals how the study of biological systems can look outward to how structures otherwise considered "external".

Contrast this with an internalist methodology. Whereas an externalist methodology takes seriously the role of external structures in off-loading and distributing the cognitive workload, an internalist methodology downplays the cognitive importance of external resources (see, e.g., McCabe, 1986). Consider, for example, the flowing patterns of rivers. Centrifugal and gravitational forces often produce common "S" shape patterns. The meandering structure constitutes the flowing rivers, but it also specifies the rivers as rivers. It therefore offers perceptual information about the rivers' identity. On an internalist view, though, the cognitive import of this fact remains elusive. Perceptual identification would be explained by referencing internal representations formed through abstraction (see McCabe, 1986, p.13). Yet such an account is overly complicated. The comprehensive and permanent structure of the rivers already stands available for perceptual use.[15]

A methodology that makes external structures central to investigation comports well with the identification of wide psychological mechanisms. When informational structures spread across internal and external parts, the search for wide mechanisms becomes much easier. If cognitive systems exploit external structures, then explanations that look to how those structures are integrated into wide mechanisms become easier to identify. As Gibson puts it: "[L]et us begin by describing the general environment of all animals and then describing the ambient environment of a signal animal at a given place and time" (1966, p.7).

2.1.2 An Alternative Metaphysics of Mind

Second, if wide explanations identify mechanisms that have component parts external to the individual, then a metaphysics that assumes cognitive states supervene only on intrinsic states fails to deliver wide explanations. Wide explanations require that cognitive states can supervene on parts outside the

[15] This mirrors Wilson's (2004, p.159) discussion of exploitative representations, which take advantage of persistent external structures in order to achieve some cognitive task at reduced effort.

individual. Wide music psychology needs to adopt a different underlying metaphysics of mind.

Such an alternative metaphysics is provided in Wilson's (2001, 2004) notion of "wide realization". Wilson (2001, p.6) points out that the internalist supervenience thesis is the result of conjoining two metaphysical assumptions: (i) the sufficiency thesis, which holds that realizers are metaphysically sufficient for the cognitive states they realize, and (ii) the constitutivity thesis, which holds that the realizers of cognitive states are exhaustively constituted by the intrinsic physical states of the individual whose states are realized (Wilson, 2001, p.4-5). For the internalist, the sufficiency thesis and constitutivity thesis hold true of the same realizers: brain states. The realizers that are metaphysically sufficient for realizing cognitive states are constituted exhaustively by the intrinsic, physical states of the individual; recall what was said about the internalist supervenience thesis in Chapter 1. However, as Wilson points out, it just isn't the case that the sufficiency and constitutivity theses always have the same realizers. The physical states metaphysically sufficient for realizing cognitive states can, sometimes, include properties or parts outside the individual. The realizers or realization base for cognitive states can be "wide".

To see how this might work, first notice that there is a distinction between the total realization and core realization of a state or property (see Shoemaker, 1981). The core realization includes only those parts that are crucial for producing or sustaining the state, while the total realization includes those core realization parts and the non-core parts that are sufficient for realizing the state. To use Wilson's (2001, p. 64) example, the core realization base of having the property of blood pressure 120/80 includes clogged arteries and a strong heart. However, these parts are not sufficient for producing blood pressure. Other non-core parts of the circulatory system are needed, e.g., capillaries, arterioles, venules, etc. It is these other parts, in addition to the appropriate background conditions such as having oxygenated blood, that form the total realization base for the property of having blood pressure of 120/80. The realizers of the property, having blood pressure of 120/80, are constituted by the physical states of the individual. The total realizers sufficient for the property of the circulatory system are entity bound; they are within the individual who bears the property.

However, this isn't always the case. Sometimes the total realization base can include external components. Consider the mechanism responsible for digestion in the giant water bug Lethocerus (see Wilson, 2010). Lethocerus injects its prey with digestive enzymes. These enzymes liquefy the prey's innards. After this process has occurred, Lethocerus then sucks up the broken down elements. The sufficient and constitutivity theses come apart here. The mechanism responsible for digestion includes components spread across internal and external parts. Even though the core parts of the realization base are contained within Lethocerus, the non-core parts, which in addition to the core parts form the total realization base sufficient for realizing the digestive mechanism, are external to Lethocerus. But if the total realization base sufficient for the digestive mechanism includes parts outside the organism, the constitutivity thesis doesn't hold. Parts internal to the organism do not exhaustively constitute the realizers sufficient for Lethocerus' digestive mechanism. Lethocerus' digestive mechanism has a wide realization base.

Adopting wide realizations as the underlying metaphysics of mind in wide music psychology offers a powerful way of conceiving of the determinative relationship between the components parts of the wide explanations. It provides a metaphysic to ground the wide mechanisms identified in wide explanations. If the mechanism identified in psychological explanations need to be, in part, externally constituted, wide realization offers one direct way for thinking about how this might work.

2.1.3 The Informational World

Finally, as we saw in Chapter 1, cognition is often taken to involve the processing of information-bearing states; states that are usually thought of as internal to the individual. However, if the methodological gaze within wide music psychology shifts outward, it isn't too far of stretch to see that the location of cognitive information-bearing structures might also change. If wide mechanisms can be formed by external structures, then any information carried by those structures becomes crucial to the investigation. Wide music psychology needs to lend an eye to the rich, dynamic informational content of the environment.

Environments are often structured in informationally rich ways (Gibson, 1966, 1986; McCabe & Balzano, 1986; Clark, 2003). From the representations of culture and language to the information carried in natural objects, environments abound in rich structures that allow them to carry, specify or afford information. A prime example is Gibson's ambient optic array. The optic array is the diffused, refracted and radiating light that fills an organism's perceptual environment. Overlaying, but not extensive with, environmental surfaces, it specifies information about the layout or structure of the environment. This is particularly evident in case of texture gradients. Texture gradients remain constant as perceivers move through their environment. As the density of optical texture increases, the scale of space is more easily defined. The equal amounts of texture represent equal amounts of terrain. In effect, texture gradients offer information about the distance of objects to the visual system. The optic array is an external information-bearing structure. Yet where internalists would view such external information-bearing structures as inputs or constraints, usually for underdetermination reasons, an externalist would view them as forming an integral part of the psychological mechanism. External information-bearing states aren't just inputs to some internal processing: They are part of the processing itself.

2.1.4 Two Kinds of Projects

In the interest of perspective, it will be helpful to take a step back and make some general comments about the kind of project being engaged in. Rowlands (1999, p.37) draws a distinction between two different kinds of projects: "psychotectonics" and "psychosemantics". On the one hand, psychotectonics concerns how to build cognitive systems. It attempts to describe how construct and explain systems as cognitive; it is a kind of engineering project. On the other hand, psychosemantics deals with how to attribute semantic content to cognitive systems. It attempts to characterize the principles by which semantic content can be attributed, e.g., belief states, to those systems deemed cognitive.

It is important to pull these two projects apart because one can be engaged in the one kind of project without necessarily being engaged in the other. Attempts to describe the attribution conditions for semantic content do not impugn on describing the design of cognitive systems that those semantic attributions apply to. This is not to say that the cognitive systems being described need not be able to support some kind of semantic attributions. But it is to say that describing how they support semantic attributions is a distinct kind of project.

To illustrate, consider the longstanding debate in the philosophy of mind over whether mental content is “wide” or “narrow”. Take the belief that water is wet. One question is what determines the content of this belief. Some have argued that it is the local environment (Burge, 1979), others that it is the belief’s relational status among other beliefs (McGinn, 1982).[16] Whatever the answer, what’s important is that the debate centres on when and why semantic attributions are warranted. It does not focus on the character of the mechanisms that enable mental states to have semantic content. It might be the case that

[16] Historically, internalists have maintained that semantic content is possessed *essentially* by internal, mental representations. For some this has meant that the semantic content must individuated relationally. In contrast, externalists have maintained that semantic content is essentially environment involving. Semantic content is not determined by its relational status among other representations but, rather, by its relation to environmental factors -- that is, how it depends on the local causal context.

mental states cannot be individuated without reference to the environment, but this does not mean that the structural basis for semantic attributions are external; they might be entirely internal. As Rowlands writes:

> If externalism about content is true and if (certain sorts of) mental states have their contents essentially, it follows that philosophical externalism about mental states is also true. However, this is perfectly compatible with the claim that an organism to which the attribution of content is warranted also possesses purely internalist states and mechanisms. (Rowlands, 1999, p.39)

There is difference between an interpretative project about the individuative criteria for mental states and an engineering project about the structural basis of mental states.

In outlining the externalist assumptions, I have described principles by which particular kinds of cognitive systems can be conceptually fashioned, e.g. wide mechanisms. Doing so, however, entails little about how semantic content should be attributed to mechanisms on the basis of the mental states they support. I have been engaged in a project of psychotectonics, not psychosemantics.

It is for this reason that I have continued to describe the vehicles of cognition as information-bearing states and not propositional syntactic structures or subsymbolic representations.[17] To do so would be to begin to encroach on the territory of psychosemantics. The externalist assumptions are more akin to design principles than individuative criteria. Explaining the structural basis of mental states within wide music psychology is not the same

[17] Thus, one benefit of this position is that it remains neutral as to whether cognition has a connectionist or classicist structure.

thing as explaining how mental states get their semantic or folk psychological interpretation. Figuring out the latter I leave to others.

2.1.5 Summary

There are three points to note with respect to the foregoing discussion. First, if there is good reason to adopt wide explanations, then there is also good reason to adopt the externalist assumptions of wide music psychology. Second, similar to the internalist tripartite assumptions, the externalist assumptions are structural constraints, drawing the theoretical boundaries of wide music psychology. Third, articulating the externalist assumptions of wide music psychology is a project of psychotectonics, not psychosemantics.

The manner in which I have arrived at these conclusions may seem odd to some. Usually one would derive explanatory commitments from theoretical principles, not the other way round. Why choose to do the reverse? It is because I take explanatory practice seriously when trying to articulate a paradigm for cognitive science (see, e.g., von Eckhardt, 1995). Instead of trying to read a philosophical framework off of scientific results, I think framework-construction is best approached taking seriously the explanations in scientific practice that are already useful and desirable. This is why I worked from explanation to framework. Whether or not this approach is viable rests or falls with how well the project more generally turns out.

2.2 Wide Research

In this section, I identify some music cognition research that can be thought of as falling under the wide music cognition rubric. I survey two types of research: adult music perception and infant melodic perception. I attempt to show that each kind of research can be accommodated within wide music psychology. The general aim is to persuade the reader that wide music psychology already has a lot going for it when it comes to research. One does not have to look far for research that exemplifies each of the externalist assumptions.

To be clear about the aim of this section, I take it that I have already established the fruitfulness of wide explanations. Thus, the discussion in this section should not be seen as an attempt to vindicate or establish wide music psychology via empirical studies. Rather, it is an attempt to identify the kinds of work that exemplify the externalist assumptions, the kind of work that music cognition should attempt to incorporate and produce more of. To use a geography metaphor, given that the conceptual terrain has been mapped, now it is time to colour in the details.

2.2.1 Adult Music Perception

The first example of wide music cognition research I want to look at is one we have already briefly encountered: Gerald Balzano's (1986) research on adult music perception. This research attempted to investigate how untrained adult music listeners perceived music, how listeners identified music qua music.

Balzano advanced the view that music perception involved the detection of objective structural "pitch-time constraints". It will be beneficial to begin with Balzano's notion of pitch-time constraints, since it is central to his account.

Consider the difference between speech and music. There are an infinite number of values for pitch and time in speech. For example, vowels remain constant across changes in pitch and time whether spoken in a high or low register. There are, however, only a relatively small number of pitch-time values in music. Furthermore, these pitch-time values stand in specific relationships. As Balzano puts it: "There is a difference in the presence of specific constraints in the global selection of pitch and time values" (Balzano, 1986, p.218). These constraints entail that there are two invariant properties generated by certain kinds of acoustic patterns. First, there are "quantal" properties, which refer to the relatively small number of values of pitch and time in music. Second, there are "generative" properties, which refer to the specific relationships that hold between pitch and time values (Balzano, 1986, p.218). For present purposes, the details of these relations are not particularly important (see Balzano, 1986, p.218 for details). What is important, though, is that the musical sounds have invariant structure: Across transformations in the melodies certain pitch-time relations remain constant.

How do these pitch-time values figure in music perception? To answer this question, Balzano (1986) had subjects listen to 38 pseudomelodies and then attempt to decide whether the sounds were more or less musical. Participants either heard pseudomelodies that used pitch pairs with identical rhythm (time structure) and different pitch or time pairs with identical pitch and different rhythm (time structure). The pseudomelodies were 70-notes in length, varied in tone duration between .14 to .475 sec, and had a two octave range between A3(220Hz) and A5(880Hz). Balzano compared the pitch-time value relationships across several conditions. In a low condition, pitch-time values had unlimited range; in a medium condition, pitch-time values were quantized but without a generative relation; and in high condition, pitch-time values were quantized with a generative relation. This experimental design meant that the high constraint conditions included quantal and generative properties, the medium conditions included only the quantal property, and the low conditions included neither.

The underlying idea was that if listeners were not responsive to the quantal and generative properties (the musical invariants), then the pseudomelodies would have been identified as musical sounding to roughly the same degree across all three-constraint conditions. However, the pitch-time constraint conditions that included the quantal and generative properties showed significantly higher effects on musical identification. When pitch-time values were quantized with generative relations, as they were in the high and medium conditions, perception of musical qualities of the pseudomelodies increased (Balzano, 1986, p.227).

These results are important for two reasons. First, they suggest that musically untrained participants resonate to higher order or invariant properties of musical sounds. Second, they suggest that musical perception does not only involve the construction of musical information from impoverished stimuli. Music perception also fundamentally involves the “pick up [of] structural constraints of the sort that distinguish music from non-music” (Balzano, 1986, p.233). These points suggest that Balzano’s research might comport well with the externalist assumptions.

First, notice the outward looking character of Balzano’s own characterization of his study: “I treat the structures being described [invariant musical properties] as present in stimulation, to be detected by a perceiver, and not something to be invented or constructed by a perceiver” (1986, p.230). It is because there is rich structure in the pitch sets of the pseudomelodies that Balzano is focused on how subjects respond; the quantal and generative properties are higher-order relationships between variables in the sounds over time. “What makes music a unique stimulus?...Our answer, it seems to me, must lie in the presence of certain rather specific constraints in the global selection of pitch and time values” (Balzano, 1986,p. 218). Music cognition is not only a mental construction; it is also constituted by the external structure of temporally extended sonic events.

Second, if the pseudomelodies are structured sonic events that are used during perception, then they are also external information-bearing structures. For the pitch-time array, as Balzano calls it, specifies the sonic events as musical -- this is why listeners were more likely to identify sounds containing the pitch time constraints as musical. It is because of the presence of “musical invariants" that subjects can identify music qua music. Temporally extended sonic events specify musical features because of their invariant properties.

Finally, notice how Balzano characterizes the relationship between the musical invariants of the pitch-time array and the auditory system:

> The metaphor I prefer for talking about [musical] perceiving renders is more like a process of tracking. Consider a standard pursuit-rotor task, or, perhaps better, the task of following the movement of a bird in flight. All that I must do to discharge the difficulty of these tasks effectively is to keep the tracking organ (eye, hand) reasonably close to the (moving) object of the tracking. (1986, p.230)

Conceptualizing the relationship between the musical invariants and the auditory system as one of constant, tight interaction is suggestive of the presence of a wide perceptual mechanism. If it is integrative coupling of internal and external components that supports the larger function, then the realization base for the psychological mechanism should also be wide. Balzano's research of adult music perception seems to already have several of the basic features necessary for wide music psychology. It therefore offers an illustrative case of how the three externalist assumptions might figure into and guide wide music cognition research.

2.2.2 Infant Melodic Perception

The second example comes from Sandra Trehub's work on infant melodic perception (1977, 1984, 1987). Trehub focused on how infants perceive musical stimuli (short melodies) as either familiar or novel. In a series of studies, Trehub and others demonstrated that infant perception might involve use of global information (see Chang and Trehub, 1977). This led Trehub et al. (1984) to conduct two experiments.

The basic experimental design of Trehub et al.'s investigation involved presenting infants with a standard or original six-tone melody and then evaluating how well they discriminated between different types of transformations of the original melody (using a head turning paradigm). Each experiment began with a training phase. Infants were conditioned to respond to the standard or original melody. After this, four experimental conditions, each involving transformation melodies, along with control conditions, were tested and compared. In the first experiment, four conditions were used. These included (i) a transposition condition, where the absolute frequencies of individual tones were changed but the frequency ratios were preserved, (ii) an octave change/contour preserving condition, where the melodic contour was preserved but the absolute and frequency ratios were changed, (iii) an octave change/contour preserving condition, where the melodic contour was preserved by changing the octave of individual tones, and (iv) a contour violating condition, where the melodic contour was changed by changing both the octave and individual tones of the original melody (Trehub, 1984, p.822).[18]

Trehub et al. were interested in how the infants would respond to each of the experimental conditions. They wanted to know how the infants would treat the transformed melodies, i.e. whether they would discriminate them as either familiar or novel. The idea was that depending on how the transformation melodies were treated this would reveal what kind of information infants were

[18] Melodic contour refers to the overall change in pitch during a melody (sequence of tones)

sensitive to. For example, if the infants only responded to the contour-preserving transformation of the second or third conditions as novel, then this would reveal that they were responsive to precise information about successive intervals in the standard melody (Trehub et al., 1984, p.826). However, after presentation of the melodies in all of the test conditions the infants were able to discriminate all of the transformation melodies from the standard; they identified all of the transformation melodies as novel, although contour preserving ones a bit less so (Trehub et al. 1984, p.826). This result failed to shed light on which type or types of information were perceptually crucial. Thus, a second experiment was needed.

In the second experiment, the experimental design was kept the same as the first, but a fifth test condition as added to increase the task difficulty. In this contour-violating condition, the tones of the standard melody were temporally rearranged. The purpose of this addition was to allow Trehub et al. to more easily identify the role of contour violation independent of changes in frequency range and octave transformation (see Trehub et al., 1984, p.826). This time the infants were unable to discriminate the transposition or contour-preserving transformations from the standard melody. The infants heard the transformation melody as the same.

Trehub et al. interpreted these results to mean that melodic contour information and overall frequency range were crucial to infants for melodic perception. “[I]nfants treat new melodies or tone sequences as familiar if these sequences have the same melodic contour and frequency range as a previously heard sequence and as novel if either the contour or range differs" (Trehub et al., 1984, p.828).

For present purposes, what's particularly interesting is that these results are quite close to those of Balzano's. Melodic contour and frequency range seem to be alternate ways of describing the pitch-time constraints articulated by Balzano. The melodic contour describes changes in pitch values over time; the frequency range describes the specific relationships between the pitches within melody. Trehub et al.'s results seem to reinforce Balzano's conclusion that there are certain acoustic patterns that contain musical invariants, structural relationships that remain constant under transformation, which perceptual systems directly interact with.

However, one might wonder whether the musical invariants aren't already encoded in the perceiving subject. If so, hasn't the external, informational structure been overemphasised? After all, Trehub et al. do talk like this at times (e.g., Trehub et al., 1984, p.827). One reason for thinking otherwise is that such a supposition is unnecessary. If the standard and transformation melodies share the same pitch-time values, it is redundant to suppose these structural features also need to be encoded by the infants' perceptual systems. Why suppose that the information has to exist within the perceiving subject if it is already present in the melodies? There seems little reason to posit costly encoding processes when the perceptual system can simply exploit constant, relational properties that already exist within the acoustic patterns. As Andy Clark aptly puts it with his 007 principle: "[C]reatures will neither store nor process information in costly ways when they can use the structure of the environment and their operations upon it as a convenient stand in for the information processing operations" (1989, p.63).[19] Furthermore, given the presence of pitch-time constraints so early in development, the central role of musical invariants in music perception seems clear. The question now is how this research relates to wide music psychology.

First, consider Trehub's interest in finding the stable patterns of melodic perception. Note how the experimental design was set up so as to reveal the presence of global, relational properties in transformed melodies. Such sensitivity to the structure of sonic patterns demonstrates an allegiance to a world-oriented methodology. For it attempts to, if not explicitly then implicitly, identify the persistent structures exploited by the perceptual system. This is what led Trehub et al. to claim that, "infants' perception of melodies can be said to be holistic or structured, with the global properties of contour and range perceived across transformations of specific properties, such as interval size and absolute frequency" (Trehub et al., 1984, p.828). Trehub's investigation demonstrates how worldly invariants can be discovered when focus is shifted from how subjects encode and construct information to how subjects resonate or detect environmental structures.

Second, notice how the musical invariants support the perceptual mechanism's function, i.e. identifying melodies as familiar or novel. The musical invariants are integratively coupled with the auditory system. They are

[19] Or, more idiomatically: "[K]now only as much as you need to know to get the job done" (Clark, 1989, p.63).

crucial to the performance of the cognitive function (melodic identification). Thus, similar to the enzymes in Lethocerus' extended digestive system (see section 2.1.2), the musical invariants are external elements sufficient for realizing the mechanism responsible for melodic perception. They are the non-core parts of the total realization base that form a wide realization basis for the underlying perceptual mechanism.

Finally, the fact that there are global, relational properties present across transformations of the original melodies is suggestive of how Trehub's work subscribes to an external information-bearing states view. For it signals how Trehub's work includes a space for conceiving of sonic events as rich, informational structures. This harkens back to Gibson's talk of how environments offer information by specification (1983, p.243). Trehub et al. not only identify structures within sound patterns that specify melodies, they show how infants detect and exploit the presence of such informational invariants for melodic perception. This is a prime example of endorsing an external information-bearing states view. Given its fit with each of the externalist assumptions, Trehub's research seems to aptly exemplify what research within wide music psychology might look like.[20]

[20] It is true that Trehub does not see her work in the way I have described. Rather, she situates herself within a tradition closer to the work of Krumhansl. But the contrast is not as great as it might seem between this research and Balzano's. For, as we have seen, each approaches the study of musical patterns under the assumption that they are complex, dynamic structural entities. In this way, they both buck the underdetermination considerations we saw that characterized research like Krumhansl's.

2.3 Wide Music Psychology and Extended Cognition

In the previous section, two types of music cognition research were described at length. The purpose of this was to flesh out the character of research within wide music psychology. Although the picture of wide music psychology is becoming clearer, I want to bring it further into focus by considering one of its more interesting (radical?) consequences: extended music cognition.

What does it mean to say that cognition is extended? It means that at least some cognitive processes literally extend beyond the boundary of the individual and into the surrounding environment; that cognitive processes are at least in part constituted by elements outside the individual. In the context of music cognition, the claim is that the cognitive processes responsible for music perception include constituents external to the cognizing subject. In this section, I argue that under wide music psychology music cognition should be thought of as an extended phenomenon.

2.3.1 Music Cognition and Wide Computationalism

At least since the 1950s the cognitive sciences have been committed to the view that cognition is a form of computation. The view has been expressed in various ways. Paul Thagard has said that, "[t]he central hypothesis of cognitive science is that thinking can best be understood in terms of representational structures in the mind and computational procedures that operate on those structures" (2010, p.6). Jerry Fodor has claimed that, "quite independent of one's assumptions about the details of psychological theories of cognition, their general structure presupposes underlying computational processes"

(1975, p.28). These authors, like many others (Pylyshyn, 1984), have held that cognition is best explained by reference to a computational vocabulary. Of course, the form of computational cognition is a matter of dispute. Nonetheless, the assumption is quite pervasive.

How has this affected the study of music cognition? It has meant that one dimension of research involves giving computational accounts of cognitive systems (see, e.g., Temperlery, 2001; Hammanaka, Hirata, & Tojo, 2006). But what if the concept of computation itself turned out to be amendable to wide characterization? What if computational systems extend beyond the individual because they are locationally wide? The promise is that if computational systems are wide, then cognitive systems are extended. It is this position that I want to develop in what follows.

Consider, then, Wilson's (1994b, 1995, 2004) notion of "wide computationalism". The basic idea is quite straightforward, and has been acknowledged, in principle, by several authors (Piccinini & Scarantino, 2011; Segal, 1989, 1997). Computational systems are wide just in case some of their computational units are not wholly instantiated "in-the-head". Wilson writes:

> [W]hy think that the skull constitutes a magic boundary beyond which true computation ends and mere causation begins? Given that we are creatures embedded in informationally rich and complex environments, the computation that occur inside the head are important part but are not exhaustive of the corresponding computational system. (2004, p.165)

Wilson here is emphasising the location neutrality of computational descriptions of cognitive processes. If the method of computational analysis is locationally silent, then computational systems can be wide or extended. Nothing about computational individuation requires computational systems only be instantiated in the head.

To illustrate, consider one of Wilson's examples: The multiple spatial channel theory of form perception of Sekuler & Blake (1990). According to this theory, an organism's visual system has channels that decompose any visual scene into four parameters: orientation, spatial frequency, contrast, and spatial phase. "On this conception of form perception, part of the task of the perceptual psychologist is to identify formal primitives that adequately describe the visual environment" (Wilson, 1994b, p.363). The multiple spatial channel theory can be thought of as involving a wide computational system because the inputs to the computational processes involve environmental elements. But even granting the prima facie coherence of Wilson's example, the question is how to identify wide computational systems, and, more pertinently, whether it can be determined that music cognition implements such a system.

How is it that one can know when relations between an organism and environment qualify as computations? Consider how this is normally done for an internal psychological process. To computationally model some in-the-head process, one first must identify and formalize the relevant primitive states. This allows the process to broken down into component features. Then, one must describe the changes between the states in terms of transition rules; they must be given a function-theoretic account (Wilson, 2004, p.167). For example, in

Marr's (1982) theory of vision the retinal image and the internal 3-D representation of the environment are the relevant input and output states, while the transition rules are computational processes that govern the transformations of the information contained within the retinal image to the 3-D representation. In sum, identifying wide computational systems requires: (i) identifying and formalizing specific properties of the environment that an organism is sensitive to, (ii) decomposing natural scenes into the parameters set out by the formal primitives, and (iii) specifying the algorithms or rules that apply to the identified primitives governing an organism's behaviour. The important point to note is that where the informational primitives of the system could be internal, they could also be environmental.

Consider, then, what happens when musical invariants are taken as the primitive states of the computational analysis for music perception. First, the musical invariants can be identified and formalized as computational units given that they are external information-bearing states. This means they can supply the inputs to the computational processes, requirement (i). Second, as was seen in both of the two research examples, the auditory system is directly responsive to musical sounds; it responds as a function of the presence of external music information structures, e.g., Balzano's pitch-time constraints. This means that musical events or scenes can be decomposed into the invariant constraints, requirement (ii). And finally, there seems to be a lawful causal relationship between the physical structures of sounds, i.e., musical invariants, and the stimulation of the auditory system, requirement (iii).[21] It seems that all three of the above requirements for computational characterization are met.

Given this, it seems that music perception can be said to involve computational processes that range across environmental and in-the-head elements. There is good reason, then, to consider music perception as part of an extended computational system. The point can also be put in functional terms: the musical invariants are used by downstream parts of the system such that they become integrated into, incorporated as parts of, a wide computational system. Call this the extended computational view of music cognition (ECMC).

However, more needs to be said about why the in-the-head plus musical invariants can be given a computational characterization. There are two reasons to think this might be the case. First, the invariants persist through time; second, they have structure in virtue of which they carry information. As Piccinini & Scarantino (2011, p.463) point out, as long as a medium or vehicle has persisting informational structure it can be given a computational characterization.

Consider again what Gibson (1966) says about texture gradients. Light is constantly diffused and reflected throughout the environment. Because of this, there are textures overlaying surfaces. As perceivers move through environments, the amount of texture corresponds to the amount of terrain. In

[21] It's worth mentioning that describing this causal relationship is not an easy task. And it has not been attempted here for good reason, because it is part of the job of cognitive psychologists. All I can do gesture at its plausibly. Nonetheless, given that the direct responsiveness of subjects' auditory system to the presence of musical invariants, it does seem plausible that such a lawful relationship exists and that it could be described in more refined terms.

effect, this means that as the density of optical texture increases the scale of the space is specified or revealed. Optical texture provides information about the environment. Analogously, musical invariants such as pitch-time constraints have order and equivalence relations that remain constant through flux or transformations (see Balzano, 1986, p.227 or Trehub, 1984, p.828); they have persistent structure. Like optical texture, they provide information to the perceiver but do not depend on the observer for they existence. Rather, perceivers stand in specific relationships to the invariants. Perceivers, as we have seen, resonate or detect the information contained or carried by the invariants persistent structure. For these reasons, musical invariants seem apt to be included within computational analysis: They are the right kind of external information-bearing states.

2.3.2 Wide Mechanisms, Computationalism, and Explanations

What of wide mechanisms? How do wide mechanisms fit with wide computationalism? Two points are suggestive of the connection. First, wide mechanisms describe the causally integrated elements that support psychological processes: some internal, some external. Second, computational modeling requires treating causal relationships as inferential processes. Taken together, these points suggest that wide computational systems are the information processing gloss on the causally operative components of wide mechanisms.

Consider that at a causal level of analysis, wide mechanisms describe the relationship between internal and external components as they serve some psychological function. But at the computational level, they can also be seen as solving some information processing problem. At this second level, wide computational systems are the abstracted information processing way of describing the underlying causal processes of wide mechanisms. Wide mechanisms are the causal bases that carry out the computational processes; they are the lower level causally integrated systems that are computationally characterized. Depending on the level of analysis, the causally operative components can either be seen as the causal substructural basis for some psychological process or the computational device that processes information. The connection, then, is one of interrelated levels of description.

This fact has implications for explanation. Wide computational accounts become, in effect, a species of wide explanations. Though one can talk about the causal mechanisms that support some cognitive function, it is sometimes explanatorily useful to move to an information processing level description (see Marr, 1982; Dawson, 1998). Wide explanations track wide mechanisms and wide computational systems track wide mechanisms just at the computational level of analysis.

2.3.3 Defending Extended Music Cognition

Here is the argument that has been developed:

(i) Cognition is a species of computation.

(ii) When the elements of a computational system include parts of the organism and parts of the environment, the computational system literally extends beyond the individual and into the world.

(iii) Music perception involves the detection of musical invariants.

(iv) The interaction between the auditory system and musical invariants is characterizable as a wide computational system.

(v) Therefore, music perception is extended.

How might this argument be challenged? It seems fair to say that the argument cannot be challenged at premise (i), since to question it there would be to reject one of the defining assumptions of cognitive science. It also seems safe to say that premise (iii) is acceptable, since, as far as I know, no one has explicitly contested Trehub's or Balzano's research.[22] However, premises (ii) and (iv) seem more contestable, since each makes controversial claims about computational cognitive systems, either in terms of the scope of the concept or the phenomena that fall under it. Thus, in what follows, I defend the argument for the ECMC at these two points.

First, Segal (1997) has suggested that the very idea of wide computationalism is confused. Wide computationalism, he argues, fails to specify whether the inputs to its functions are environmental properties or internal representations. Because of this, extended music perception is predicated on a convoluted notion.

In his original presentation of wide computationalism, Wilson claimed that spatial navigation, what's known as dead reckoning in some animals, could be given a wide computational account. He claimed that spatial navigation in ants takes as its inputs "solar heading, forward speed, and a representation of the

[22] This is not to say that someone could not still question the argument on these grounds.

solar azimuth" and delivers as its output "a representation of the creature's position relative to some landmark" (1994b, p.366). However, as Segal (1997, p.34) points out, whereas the first of these two features, solar heading and forward speed, are properties of the environment, the third, representation of solar azimuth, is a property of the organism. It is therefore unclear how properties of both the environment and internal representations can both be used in the computations of a system; surely it can only be one. At its core wide computationalism trades on an equivocation.

This concern is not particularly pernicious. It can be conceded that computations cannot involve both representations and environmental properties and still maintained that in the present context music perception computations take as their inputs only environmental elements. This is because the music invariants which music perception takes as its computational input have persistent structure under transformation. The music invariants are wholly external information-bearing structures. Characterization of the musical invariants as inputs to a wide computational system requires no further reference to internal representations. The musical invariants carry the requisite musical information.[23] The worry that wide computationalism problematically slides between inner and outer information-bearing states ceases when the musical information serves as the input to the extended music perceptual processing.

[23] To give credit where credit is due, Wilson hints at this response when he says: "Once we take this step the interaction between the information processing structure inside organisms and information bearing states outside of them becomes central to a computational account" (2004, p.171).

Second, one might worry that while it may be the case that cognition is computation, it is certainly not the case that all computation is cognition. Since calling something computational is just one way of interpreting it (see Stufflebeam, 1999), it is possible to describe almost anything, including music perception, as computational. However, this does not mean that such computational characterizations make things part of an extended cognitive system.

This worry trades on a conflation between computational explanation and computational modeling. The question is not whether physical systems can be given computational descriptions, but whether the physical systems themselves can be thought as performing computations. Finding out whether a physical system can be formalized and captured by transition rules is a difficult task, but it is precisely part of the task of research within wide music psychology. Discussion of Balzano's research was meant to show that important aspects of the acoustic environment could be formally computationally analyzed. The invariants deliver musical information to the auditory system such that the physical components of the environment plus in-the-head components can be thought of as performing computations. The present argument provides evidence for viewing music perception as a computational explanation of a wide computational system, not merely a target for computational modeling (see Piccinini & Scarantino, 2011, p.6 or Wilson, 2004, p.168).

2.3.4 Conclusion

In this chapter I have provided a road map for future wide music cognition research. I have shown both what elements are needed and where to look. By outlining the three externalist assumptions and identifying and describing two types of research, I have mapped out the conceptual and empirical terrain of wide music psychology. What's more, I have shown how going wide about music psychology has the consequence of extended music cognition. Given all this, in the next chapter I want get a bit more specific and situate the ECMC with respect to two other views that have been offered under the umbrella of "extended" music cognition.

Chapter 3 - Extended Music Cognition

3.0 Chapter Overview

There are two goals of this third chapter: (i) to situate the ECMC with respect to two other extended views of music cognition and (ii) to evaluate how each view stands up to three internalist challenges. I begin by outlining the two alternative approaches and then turn to how each relates to the ECMC. The purpose of this discussion is to chart the different ways one might approach cognitive extension within wide music psychology.

3.1 Music Cognition and Cognitive Extension

In Chapter 2 I argued that the constituents of cognition extend beyond the individual: because the constituent elements of the computational processes involved in music cognition included invariants outside the individual, music cognition was part of an extended computational system. Yet while the notion of wide computationalism provides the basis for one direct argument for cognitive extension, it is not the only such argument. One finds at least two other major approaches to cognitive extension within the extended cognition literature.

First, there are "coupling arguments". Arguments of this stripe attempt to establish cognitive extension by reflecting on how external processes or components can become reliably causally connect to internal processes such that they form a composite system with functional gain. The claim is that when coupling conditions are right, processes become extended (see, e.g., Clark, 2008). Take, for example, mathematical problem solving. During mathematical problem solving the numbers on a page are tightly integrated with internal, neural resources via action/perception cycles involving pen and paper. The person-plus-body-paper system has functional capabilities otherwise not available to the lone problem solver. Under such conditions, it makes sense to treat the composite system as a cognitive system in its own right. The external vehicles of pen and paper, when coupled to the problem-solving agent, extend the problem solving process.

Second, there are "parity arguments". These arguments motivate cognitive extension by appealing to functional equivalence. They claim that insofar as external elements play the same functional role as internal elements in support of some cognitive function, they can extend cognition. So, for example, if the information in a notebook is deployed in a functionally equivalent manner to how it would have been used had it been stored in internal, biological memory, then that external information can form the constitutive base for an extended memory. And, since the memory retrieval process trades in an extended cognitive state, it too becomes extended (see, e.g., Clark & Chalmers, 1998, p. 13).[24] In what follows, I examine two accounts of music cognition that adopt these two general approaches to cognitive extension.

[24] Parity arguments rely on what Clark calls the parity principle: "If, as we confront some task, a part of the world functions as a process which, were it done in the head, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world is part of the cognitive process" (Clark, 2008, p.222).

3.1.1 The Musically Extended Mind

The first view to consider is Joel Krueger's (2014) "musically extended mind". Krueger's thesis is that "music serves as an external (i.e., outside-the-head) resource that can profoundly augment, and ultimately extend, certain endogenous capacities" (original emphasis) (2014, p.4). He offers the following line of reasoning in support of this claim:

> [A]s with the performer-instrument relation, listener and music similarly form an integrated system…Within this system, the listener uses music (via musical affordances) as a kind of "aesthetic technology" for regulating and transforming their behaviour, attention and emotion. In this sense, then, music can – when integrated with an appropriately responsive listener – function as an external cognitive and affective resource. Via processes of synchronization and bodily entrainment, music takes over some of the regulatory functions…[functions] which the listener exploits to access a more nuanced means of emotional refinement. (Krueger, 2014, p.7-8)

Here is a plausible reconstruction of Krueger's argument:

(i) Music offers particular cognitive and motor engagements (i.e., musical affordances).

(ii) Listeners become entrained to musical affordances such that they form an integrated system.

(iii) The integrated music-listener system actively drives and augments various emotional and attentional regulatory processes.

(iv) Therefore, music -- or more specifically, musical affordances -- extends emotional and attentional regulatory processes.

(v) Therefore, music expands and extends novel emotional experiences.

Krueger adduces several studies in favour of this argument. Briefly consider two. First, Krueger points out that music often involuntary elicits patterned motor responses from children in the form of swaying or bouncing behaviour, particularly for buoyant melodies (2014, p.4). Such evidence supports premise (ii). It demonstrates music listeners' ability to synchronize or become entrained to musical structures via motor engagements. Second, Krueger notes that consonant music augments infant development by guiding regulatory competences such as respiration (2014, p.7). This kind of evidence supports premise (iii). It shows that music decentralizes cognitive control and offloads regulatory processes onto external structures for functional gain.

There are two more pertinent points to note about Krueger's argument and the account it sustains. First, there are two targets of cognitive extension in Krueger's account: (i) there are emotional experiences and (ii) there are subconscious regulatory processes. Since discussion up to this point has largely dealt with subconscious cognition, I will continue this trend and focus only the regulatory processes. This will also help to simplify discussion. The more basic regulatory processes are more interesting, since they are the ones that are directly integrated with the musical affordances via motor engagements.[25]

Second, notice that Krueger offers a coupling argument. Recall that at root coupling arguments deliver cognitive extension through reflections on

[25] The extension of the novel emotional states only comes after the regulatory processes form an extended system.

tight, reciprocal causal connections. Krueger's account adopts precisely this argumentative strategy. Premise (i) establishes the causal connection (i.e., coupling) between listeners and music via the notion of entrainment. Premise (ii) then characterizes the coupled music-listener relationship in terms of an integrated system. Add to this premise (iii)'s further elaboration of the kind of coupled system that is formed (i.e., one with functional gain). And, the cognitive extension claim in (iv) should come as little surprise. As we can see, for Krueger, cognitive extension follows on the heels of functionally gainful causal coupling that occurs between musical affordances and listeners' active motor engagements.

To illustrate the point better, consider an analogous coupling argument about gestures. Drawing on Goldin-Meadow's research on gestures, Clark (2008) points out when gesturing is allowed during spatial reasoning, individuals decease the difficulty of the task (for details see Goldwin-Meadow, 2003). This signals that gestures do much more than facilitate communication. They also offload or distribute cognitive work. For Clark, this signals that they form an integrated coupled system with the internal, neuronal processes; one that has functional capabilities not otherwise possessed by neuronal processes alone (i.e., better spatial reasoning abilities). According to Clark, within this integratively coupled system the gestures are not just causal influencers they are constitutive. The cognitive processes involved in spatial reasoning are extended through the cognitive system of which gestures are a constitutive element.

The parallel to Krueger's account is quite clear. Because listeners are causally coupled to musical affordances via bodily engagements -- producing functional gain -- they form an integrated system that can be treated as a cognitive system in its own right. In both cases, it is a specific form of causal coupling (i.e., functionally gainful causal coupling) that results in cognitively extended processes. Much like the neuronal-gestural system, the music-body-neural system sustains and enhances certain attentional, motor and emotional regulatory processes such that it extends those processes. As Krueger writes:

> When we engage in bouts of musicking [active engagement with music], we potentially use music to become part of an integrated brain-body-music system – and within this extended system, musical affordances provide resources and feedback that loop back onto us and, in so doing, enhance the functional complexity various motor, attentional and regulatory capacities. (my italics) (2014, p.4)

3.1.2 Expression and Extended Cognition

The second view to consider is Tom Cochrane's (2008) "expression and extended cognition". Like Krueger's view, Cochrane's account takes its lead from the tight, causal interaction that obtains between cognizers and musical structures. However, unlike Krueger's account, Cochrane is focused not so much on how musical affordances interact with regulatory processes but, rather, with how music interacts with the cognitive processes responsible for musical expression. Cochrane's general thesis is that musical expression is, in part, constituted by external, musical structures.

Cochrane's general thesis, however, breaks down into two more specific ones: (i) that during performances instruments extend musical creation and (ii) that music extends emotional states.[26] According to the first claim, instruments extend musical creation during improvisational jazz. According to the second claim, musical patterns extend emotional states during jazz performances. Furthermore, the second claim breaks down further into a stronger and weaker version. The weaker version claims that music only elaborates internal emotions; the stronger version claims that it replaces those emotions. For present purposes, this distinction is not particularly important, because both versions make cognitive extension claims. The only difference is that each version varies on the extent to which the extension holds. The weaker version claims that emotions are extended partially, the stronger that they are extended completely.

Cochrane deploys two arguments to establish claims (i) and (ii). In what follows, each argument is presented in its barest form.[27] First, with respect to claim (i), Cochrane (2008, p.333) argues:

(i) If we consider the simple task of generating notes, it is clear that the musician's interaction with his instrument allows him to do this…[in] the particular way in which the notes are formed is a matter

[26] By "emotion" Cochrane means emotional state, not conscious experience of an emotion. He writes: "Note that in neither case do I identify the emotion with the conscious experience of the emotion. Several theorists argue for the possibility of unconscious emotions. So my claim here is only that the emotion is partially constituted by the music, not the experience of the emotion" (2008, p.330). In this way, Cochrane's argument comports well with my continued focus on subconscious cognition.

[27] There are more bells and whistles to the arguments as Cochrane deploys them, but these seem to be the crucial elements that each rises or falls with.

sensitively responding to the capacities and affordances of the physical object.

(ii) The musician and his instrument form a single tightly coupled system.

(iii) The instrument itself helps to decide the character of each note the musician then endorses.

(iv) Thus, when competing the cognitive task of choosing what exact notes to play [musical creation], the instrument is part of an extended loop between the musician's brain, the muscles in his hands or lips and the keys of the instrument.

Second, with respect to claim (ii), Cochrane (2008, p.335-337) argues:

a) Bodily changes are what constitute emotional states. (p.335)

b) By directly modifying the bodily pattern and monitoring the results, we can experiment with the emotional state. (active manipulation – control) (p.335)

c) Thus, expression can allow a more sensitive response to one's situation. (functional gain) (p.335)

d) The patterns in the music effectively play the same role as his inner bodily changes in relating to the overall bodily pattern…because they do vital work in generating and maintaining the emotional content. (p.336)

e) So, equally we should include the music as constitutive part of the emotion. (p.337)

As can be seen, the first argument is based on coupling considerations. It has all of the essential features. First, it claims that there is a tight, causal connection between the musician and instrument, premise (i); second, it claims that the causal interaction results in a "coupled system" (between instrument and musician), premise (ii); and third, it claims that the coupled system is responsible for some enhancement to the performance, premise (iii). However, since discussion has already been given to this style of argument, it will be more instruceducative to look at Cochrane's second argument.

The second argument claims that extension is delivered via "functional equivalence". This is a prime example of a parity-style argument. Why? Because it claims that if musical patterns in the world play the same functional role as bodily patterns, then they can qualify as constituents of emotional states. Emotional states are externally constituted because musical patterns can be substituted for bodily patterns as the underlying realization base of the states.

Adams and Aizawa (2008) point out that this style of argument, which they call the cognitive equivalence argument, has a major and a minor premise: "[the] major premise strikes us a something like a logical or conceptual truth, namely, that any process that is cognitively equivalent to a cognitive process is itself a cognitive process. The minor premise maintains that this or that processes spanning the brain, body and perhaps environment is cognitively equivalent to a cognitive process" (2008, p.133).

Where are these major and minor premises in Cochrane's argument? Somewhat oddly, they both reside in premise (d). This oddity has to do with the fact that Cochrane's argument has an extra step. It claims that the functional equivalence is not between brain processes and environmental processes, but between bodily processes and environmental processes. The reason for this difference is that at the outset of his account Cochrane assumes that emotional states are constituted by bodily patterns. He writes: "I will assume that emotions are essentially constituted by patterns of bodily changes. These patterns of bodily changes are registered in the brain, which then generates the felt experience of the emotion" (2008, p.229). While this assumption has a lot packed into it, the important point to note is its implication for the parity argument; namely, that in terms of the larger functional equivalence Cochrane is trying to establish, it entails that the equivalence relation holds between bodily patterns, established in premise (a), and the external, music patterns, established in premise (d).

Note also that premise (d) can be cashed out to a greater or lesser degree. As Cochrane argues, music could augment internal vehicles which otherwise sustain emotional states or it could completely replace the need for such internal states. Either way, though, the external musical patterns gain constitution status in virtue of their functional equivalence. What determines the underlying realization base is whether the musical patterns are functionally substitutable for the bodily patterns which constitute emotions. Regardless of whether premise (d) is couched in terms of elaboration or replacement, the argument goes through to an extended conclusion, however. In sum, then, Cochrane offers a parity-based argument for cognitive extension, claiming that because musical patterns elaborate or replace bodily patterns in sustaining and generating musical expression, they extend those processes into the world.

However, before moving on, I want to flag one concern with how Cochrane's constructs his account. Early in his account Cochrane writes:

> A key part of arguments for extended cognition (and externalist views of mind generally) then involves a supervenience claim: If two identical cognitive processes or states internal to the body can nevertheless have different content due to differences external to the body, then the constitution of the mental state must partially supervene on those relevant aspects of the environment. (2008, p.332)

Notice that in the antecedent of this conditional Cochrane mentions "content", but in consequent he mentions "constitution". Here Cochrane has conflated "content externalism" with "vehicle externalism" (see Hurley, 2010). The difference is subtle but important. Content externalism refers to a view about the determinative relationship that obtains between cognitive states and the environment; vehicle externalism, though related, refers a view to the constitutive basis or underlying realization base of those contentful cognitive states.

In Chapter 2 I drew a distinction between projects that were involved in psychotectonics and those involved in psychosematnics. Psychotectonics deals with how to build cognitive systems; psychosemantics deals with how to attribute semantic content to those cognitive systems. I was careful to distinguish these two projects precisely for the reason that establishing claims about one didn't necessarily imply another about the other. Even if external vehicles determine or fix the content of cognitive states, this doesn't mean that the states themselves have to supervene on external vehicles. The realization base could still be entirely internal. In describing the motivation for extended cognition, Cochrane has equivocated between which of the two projects establishes extension. Although this is somewhat understandable, it will, as we will see, complicate his account's ability to deal with some internalist challenges. Now that the alternative views of extended music cognition have been outlined, it is time to compare them with the ECMC.

3.1.3 Coupling, Parity, the ECMC, and Grains of Analysis

What is the relationship between the extended views of music cognition just described and the position outlined in Chapter 2? How do Krueger's and Cochrane's accounts compare with the ECMC? The answer, I submit, is that they relate in terms of their level of analysis. To begin, consider David Marr's tri-level typology of cognitive systems:

> At one extreme, the top level, is the abstract computational theory of the device, in which the performance of the device is characterized as a mapping from one kind of information to another, the abstracted properties of this mapping are defined precisely, and its appropriateness and adequacy for the task are demonstrated. In the centre is the choice of representation for the input and output and algorithm to transform one into the other. At the other extreme are the details of how the algorithm and representation are realized physically. (1982, pp.24-5)

According to Marr, to fully understand cognitive systems, analysis needs to be given at three interrelated levels; Marr calls these the computational, representational and implementational levels of analysis. As Dawson describes the approach, "[t]he tri-level hypothesis offers an approach to cognitivism in which method and theory are each given an appropriate emphasis" (1998, p.290).

At the computational level, investigators ask what function a system performs; they address what the system does. At the procedural level, investigators ask how the system accomplishes its function, what algorithms it employs. And at the implementational level, investigators attempt to explain how a cognitive system materially carries out its function; investigators look for what physical constituents give rise to the cognitive system (for further details see Marr, 1982, ch.1; Dawson, 1998, ch.1). At each level of analysis, a different kind of question is asked: What the system's function is; by what procedures the system accomplishes its function; and how the system physically realizes its function.

However, as Shapiro (2000) points out, the distinction between the levels of analysis is not as tidy as Marr's typology would have it seem. This is because -- at least with respect to the computational level -- the 'what-is-its-function' question can also be asked within levels. Certain levels of analysis admit of further refinement; they are "grain-sensitive". "A healthier view of cognitive science recognizes that one investigator's computational level is another's algorithmic level; and one investigator's algorithmic level may be another's implementational level" (Shapiro, 2000, p.441).

For example, at a coarser grain of functional analysis the visual system has the function of informing perceivers about the environment. This might mean that analysis involves explaining a variety of processes such as depth perception, motion detection, color perception, shape representation, etc. However, at a finer grain of analysis, the visual system also has the function of shape representation. At this further grain of functional analysis, explanations might have to describe processes such as shape-from-shading and structure-from-motion. In both cases, the grain of analysis resides at the functional level as each grain of analysis deals with the function of the visual system. However, each grain of analysis also answers different 'what-is-its-function' questions. Both grains offer functional characterizations of the visual system. It is just that one does it at a finer and one at a coarser level.[28]

The uptake is that cognitive extension views should be understood as being pitched at "different grains of analysis". Coupling, parity, and wide computational approaches to cognitive extension each address different functional levels of organization for cognitive systems or mechanisms under different descriptive guises.

First, there is a behavioral grain, where cognitive states, and therefore systems, are functionally described in terms of the behavioral competences they support. Parity arguments are pitched at this grain, since, as we saw with Cochrane's second argument, extension is delivered via functional substitutability of cognitive states. Part of the reason for this is that the

[28] A similar point holds for the other two levels of analysis. "The implementational level descriptions might be at a neural level or, depending on the grain of analysis necessary for the investigation at hand, might descend to the molecular level. Neurons, after all, have a job to do and must be made of something that is capable of doing this job" (Shapiro, 2000, p.441).

functional characterization needs to be coarse or thin enough, given that in many other functional respects external vehicles are notably different, to establish equivalency between states.

Second, there is a causal grain. At this grain, investigation is couched in terms of how causal components are integrated in the service of some more specific function. Coupling arguments are aptly pitched at this level, because they deal with how functions are subserved by causal connections. There is a sense in which this grain can be quite higher or lower depending on the complexity of causal properties involved. This is the level of the wide mechanisms discussed in Chapters 1 and 2.

Third, there is a computational grain. Here states are described in terms of the computational processes that range over them. This kind of analysis targets external information-bearing states such as the musical invariants. Analysis is given to how computational processes support specific functions such as music identification. This grain is more akin to describing the function of state representation in the visual system, rather than to the global function of action guidance that the behavioral or causal grain might address.

Now, even though I have only given a preliminary sketch of each grain of analysis, this should suffice to show how the cognitive extension views relate. Figure 3.1 depicts roughly the typological view I have in mind.

Level of Analysis	Grain of Analysis
Functional Level	Behavioural Grain (Parity Arguments)
	Causal Grain (Coupling Arguments)
	Computational Grain (Wide Computational Arguments)

Figure 3.1. A typological representation of cognitive extension views.

Each view of cognitive extension views consists of different descriptive grains. The fit of each grain depends on the phenomenon used to achieve extension. Indeed, phenomena themselves can be pitched at different levels of description, so how the views of cognitive extension relate can be quite complicated.

To illustrate, consider Krueger's account. Krueger used musical affordances to achieve extension. In Gibsonian theory, musical affordances refer to invariants of invariants; they are the invariants offered by combinations of invariants (see Goldstein, 1981, p.193). As invariants of the invariants, affordances are a complex phenomenon. Notice, though, that Krueger's causal grain is of a higher order than the musical invariants that extend computational processes in the ECMC. In saying that musical affordances extend cognition via causal integration, Krueger adopted a causal grain of analysis. Whereas the computational grain extends simpler external information-bearing states (musical invariants), the causal grain extends informational quite complex entities (invariants of invariants). The relationship is one of ascending or descending (depending on the epistemological starting point) grains of organizational complexity within the functional level of analysis.

This makes sense given what was said about mechanistic explanations in psychology in Chapter 1. If psychology, or better yet music psychology, is in the business of identifying the mechanisms responsible for cognitive behaviour, then it stands to reason that if the mechanisms are hierarchically organized according to function, so too would the analysis. Cognitive extension comes at grains corresponding to the functional organization of the mechanism it identifies. What I have tried to do in articulating the grains of analysis framework is provide a way for thinking about how cognitive extensions views relate. It is more a heuristic tool, rather than a taxonomic apparatus. Obviously more needs to be said, but the minimal take-away should be that, like the analysis given of wide mechanisms and wide computationalism in Chapter 2, cognitive extension admits of several types of interrelated levels of analysis.

Before closing, there is one final point to make about the relationship between the different views of extended music cognition. Whereas the ECMC extends musical-cognitive processes, the coupling and parity arguments of Krueger and Cochrane extend non-musical cognitive processes. Recall, for instance, that Krueger claimed that: "music serves as an external (i.e., outside-the-head) resource that can profoundly augment, and ultimately extend, certain endogenous capacities" (2014, p.4). Here the cognitive processes that are extended are not the ones responsible for processing music. Rather, they are other, non-music cognitive processes. This point is worth mentioning, not so much because a great deal hangs on it, since it would only matter if someone had a narrow view of music cognition. Rather, I mention it because it illustrates that cognitive extension can and does apply to several dimensions of music cognition studies. It includes views that address both musical processing directly and views that address the relationship between music and other aspects of cognition.

3.2 Extended Music Cognition and Some Internalist Challenges

Having outlined and compared the extended views of music cognition, I now want to turn to evaluating their relative strengths and weakness. To do this, I will stack each extended view against three internalist challenges. The aim is to persuade the reader that the ECMC has the most theoretical power behind it.

3.2.1 Coupling, Constitution, and Extended Cognition

The first internalist challenge to consider is Adams and Aizawa's "Coupling-Constitution Fallacy" (CCF). In what has now become standard move in the extended cognition dialectic, Adams and Aizawa (2008, 2009) argue that proponents of extended cognition make an important conceptual error: They confuse the causal importance of external elements with the constitutive basis of cognitive processes. There is a distinction between "causes" and "constituents" that proponents of extended cognition fail to respect. Just because some external process or component is causally connected (coupled) to some cognitive process, doesn't thereby mean that the external process or component becomes part of the cognitive process. "[W]e cannot assume that causally coupling a process X to a cognitive process Y is sufficient to make X a cognitive process" (Adams & Aizawa, 2008, p.93). Observations establishing the first kind of relationship fail to establish the second kind.[29]

Moreover, according to Adams and Aizawa, there are two versions of CCF. First, there is a simple coupling version of the fallacy, which says that proponents of extended cognition move, either implicitly or explicitly, from observations about the causal connection (coupling) of internal and external processes to the further constitution claim that the two processes form one "extended" process. Second, there is a systems version of the fallacy, which says that proponents of extended cognition begin by identifying causally coupled components (some internal, some environmental or bodily) and then shift to talking about how those causally coupled components form an integrated (extended) whole or system, and from there move to the conclusion cognition extends into the world or body. Schematically, the CCF looks like this:

1. Y is a cognitive process
2. X is causally coupled to Y
3. X and Y form an integrated process or system
4. X is part of an extended cognitive process or system

where x = external/environmental process or component.[30] If the proponent of extended cognition moves from (1) and (2) to (4), it is a simple coupling

[29] Block (2005) and Rupert (2004) have also drawn this distinction between "causes" and "constituents". However, because Adams and Aizawa's treatment is more fleshed out and has received more attention, I will center discussion on their formulation.
[30] I borrow this formulation of the CCF from Wilson (2010, p.215).

version of the CCF. If the proponent moves from (1), (2) and (3) to (4), then it is a systems version of the fallacy. According to Adams and Aizawa, the simple version offenders include Wilson (2004, p.194), while system version offenders include van Gelder (1995, p.373) and Clark and Chalmers (1998, pp.8-9)

Support for the coupling-constitution distinction comes from a number of homely, intuitive examples. Briefly consider two (Adams & Aizawa, 2008, ch.7). First, pendulums tend to synchronize swinging patterns. Noticing this fact, however, doesn't mean that one should be moved to claim that the pendulums' swinging are constitutive parts of one another. They are clearly still two separate processes. Even if two processes are systematically coupled, this doesn't mean that they are part of each other. Second, nothing about the fact that the cooling system of an air conditioner unit is causally connected to the electrical system that powers the unit supports or licenses the claim that the two systems form a larger composite cooling-plus-power system. Expressed as purely causal-coupling claim, the observation is fine; but, expressed as a constitution claim, it fails to appreciate that while cooling is carried out in the cooling system and powering in the electrical system, nothing is carried out by or in the composite "cooling-plus-electric system". The question now is how each extended music cognition view handles this challenge.

First, recall the inference pattern Krueger offered. First, he observed that emotional/attentional processes were synchronized or coupled via entrainment with musical affordances. Then, he claimed that these internal regulatory processes and the external musical affordances formed an integrated system, one with functional gain. From this, he concluded that the emotional and attentional processes were extended. At face value, this seems pretty close to the systems version of the CCF. Consider how the argument looks recast in CCF form:

I. Emotion/attentional regulation are cognitive processes.
II. Musical affordances are causally coupled to regulatory processes.
III. Musical affordances and regulatory processes form an integrated system.
IV. Therefore, emotional/attentional regulatory processes are extended.

Krueger's argument seems to slide between observations about causal integration and constitution claims about emotional/attentional cognition via claims about integrated systems; a seemingly good example of the systems version of the CCF.[31]

Next, consider Cochrane's account. With respect to Cochrane's first argument, recall Cochrane claimed that, "the musician and his instrument form a single tightly coupled system" (2008, p.333). On this basis, Cochrane then concluded that, "when competing the cognitive task of choosing what exact notes to play [musical creation], the instrument is part of an extended loop between the musician's brain, the muscles in his hands or lips and the keys of the instrument" (2008, p.333). However, this looks like a coupling-extension

[31] If the slide occurs anywhere, it occurs between premise (III) and the conclusion in (IV). It is there that Krueger makes the crucial connection between synchronization with musical affordances and the inclusion of emotional processes as part of the resultant cognitive system.

shift. Cochrane has moved directly from the formation of an coupled musician-music system, based on observations about tight, causal connection, to an extended constitutive claim about the cognitive processes involved in the integrated system.

However, with respect to Cochrane's second argument, things are a bit more complicated. This is because whether or not the second argument commits either version of the CCF depends on how premise (d) is read. For without the functional equivalence used in (d), the argument can't go through to extension. If (d) is interpreted causally, then it could seemingly be an example of the simple version of the CCF. For it would say that because the musical patterns causally affect the bodily patterns that constitute emotions that they therefore also constitute the emotions. If, on the other hand, the relationship is supposed to be functional, then it might be a making a more sophisticated claim about constitution and function. For example, that when some functional criterion is met constitution follows. Yet, the account is equivocal. This ambiguity is, in part, created by Cochrane's conflation of content and vehicle externalism. It is unclear how the determinative relation that connects the function of the musical patterns with that of the bodily patterns should be understood; whether the determinative relation Cochrane relies on in premise (d) should be thought of as causal or constitutive. Given this, it might be best to leave the evaluation blank for now.[32]

[32] Though, in fairness, the most likely interpretation is that Cochrane means the determinative relation is causal and therefore subject to the system's version of the CCF.

Finally, consider the ECMC. Does the ECMC pay sufficient attention to the coupling-constitution distinction? The main reason for thinking that it does is the direction of argumentation that was originally offered for the ECMC. The argument for the ECMC flowed through computationalism, not causation. It was only because the external information-bearing states and internal processing were capable of being computationally explained in terms of a wide system that the cognitive extension followed. The view assumed that cognition just was computation and then asked where that cognition might be occurring. There was no inference directly from the causal coupling of musical invariants to the auditory system to extended cognition. Extended cognition followed merely in virtue of the computational characterization that was amendable to the invariant-auditory system conjunct. In light of this, the ECMC seems to subscribe to neither version of the CCF. Thus, there seems to be good reason to suppose that even granting the coherence of the CCF, which many authors have not (see Shapiro, 2011; Clark and Wilson, 2009), the ECMC avoids the challenge.

3.2.2 Motley Crew and Extended Cognition

The second internalist challenge to consider is what Shapiro (2011, p.189) calls the “motley crew problem”. This challenge claims that theories of extended cognition fail to be scientifically tractable. Wilson (2002), for example, has charged extended cognitive systems with failing to be robust enough so as to be amendable to scientific investigation. “If we recall that the goal of science is to find underlying principles and regularities, rather than to explain specific events, then the facultative [temporary] nature of distributed cognition becomes a problem" (Wilson, 2002, p.631). Rupert (2004) has claimed that external processes differ so significantly from internal ones that they should be treated as theoretically distinct. Talking about extended memory he writes, "the external portion of extended memory states (processes) differ so greatly from internal memories (the process of remembering that they should be treated as distinct" (2004, p.19).[33]

At the centre of these worries is a general pessimism about the potential of composite individual-world processes or systems to be scientifically accessible. Shapiro gives the problem the following form:

(i) There are cognitive processes wholly in the brain that are distinct from processes crossing the bounds of the brain.

(ii) Processes within the brain are well-defined, in the sense that they can be the subject of scientific investigation, whereas processes that cross the bounds of the brain are not well-defined -- are a motley crew -- and so cannot be an object of scientific investigation.

(iii) Therefore, cognitive science should limit its investigations to processes within the brain.[34]

Two points buttress the argument. First, premise (i) is plausibly supported by the fact that neuronal, chemical, and electrical processes in the brain are, in fact, different from those that occur outside the brain. Second, premise (ii) is supported by the fact that when external props or resources are made to

[33] Rupert (2004, p.391) calls this concern the problem of cognitive bloat.
[34] This formulation of the problem is taken from Shapiro (2011, p.189).

perform cognitive functions, they do so in ways very different from internal (neuronal) resources.

If external processes and systems operate on principles entirely distinct from those governing the brain, then there is little chance of finding a unifying framework capable of integrating both into a coherent whole, so the worry says.[35] The portions of the world to which the brain might be connected constitute a "motley-crew". They are collections of largely unrelated, distinct elements. It seems unlikely that they will be brought under a common framework that also includes internal brain-bound processes. Thus, holding that external vehicles extend cognition into the world undercuts science's ability to uncover the principles and regularities that make it a profitable endeavour. At best, an extended cognitive science would only allow insights into context-specific cognitive processes or systems. The question is what the prospects of an individual-world music science according to each view are.

To begin, consider what an extended musical science for Krueger might look like. As we saw, Krueger's account made use of Gibson's (1969, 1986) notion of affordances. Affordances are invariant combinations of invariants, higher order invariants that specify what invariant structures offer or afford to organisms. According to Krueger's account, musical affordances are the external vehicles that are causally integrated with the music-listener so as to form and extended cognitive system. It is through the notion of musical affordances that an answer to the motley crew problem can be given. For

[35] Adams and Aizawa (2008) gloss this point in terms "marks of the cognitive". The general thought is that internal cognition has certain unique causal regularities that are not shared by external processes or elements and so external elements cannot be studied in the same way.

affordances already have a home within the rich explanatory framework of ecological psychology.[36] Krueger takes advantage of this usage when looking to construct his extended music science. The ecological laws of acoustics, and more pertinently music, offer one route to piece together the perceiver-environment (affordance) interactions. Now, no such project currently exists. Nonetheless, insofar as the motley crew problem is pessimistic about the chances of a common framework unifying the physically diverse external vehicles, the ecological study of music offers some reason to be optimistic about Krueger's chances.

However, one potential drawback of adopting this line of response is that it still leaves open the burden of explaining how music-listeners detect affordances.[37] As things stands, Krueger's account leaves something to be desired. To use what individuals "synchronize with" to define what musical affordances people perceive is somewhat circular. Such a relational definition doesn't really explain why music should offer the affordances it does; why it is it is perceived as affording certain kinds of motor and cognitive engagements. As Goldstein (1981, p.193) pointed out, ecological psychology has enough trouble specifying ordinary invariants. If affordances are supposed to be higher order invariants, providing an informative account of them is going to be an even bigger challenge.

[36] Ecological psychology approaches cognition as a phenomenon that emerges out of continuous environment-organism interactions. It is closely associated with the work of J.J. Gibson (1966, 1986).
[37] This was part of Fodor and Pylyshyn (1981) original criticism of Gibson's ecological account of psychology.

Cochrane's chances look even worse. Cochrane's account said that musical patterns could form the realization base of emotional states, because they played the same functional role as bodily states. However, Cochrane's account offers no way to address what the relevant features of musical patterns that enable extension are. For example, Cochrane says that musical patterns resemble the "dynamic qualities" of emotional states. This might be true. But as far as systematically relating them, other than by specifying the functional roles each might play, the music psychologist is left in the dark. There is seemingly no way to connect the internal regularities with external ones other than through gross functional equivalences. As far as framework for unification is concerned, functionalism offers little help.

The wide computational framework of the ECMC offers the best response to the motley crew problem. By retaining the central conception of cognition as computation, the ECMC has the resources to address the unruliness of the musical environment. While the motley of music patterns may be a difficulty from the perspective of the auditory sciences, from the perspective of computationalism, they are quite unified. As I tried to show in my discussion of the ECMC, it is not just that musical invariants are computationally describable; they are capable of being given computational explanation. It is because the musical invariants already have information-bearing structure that they can be brought under the common vocabulary of computational analysis. The impartiality of computational analysis means that adopting this approach provides the theoretical resources necessary to investigate brain-body-world hybrids. As should be clear, in as much as internal, brain processes can be brought under computational description, so too can external ones. The computational modeling of extended cognitive processes can carry on in similar fashion to that of an internalist computational modeling.

3.2.3 Scientific Conservatism and Extended Cognition

The final challenge to consider is what I will call the "argument from scientific conservatism" (ASC). Robert Rupert offers the best example of the challenge. He writes:

> We can grant that cognition often involves intimate interaction with its environment…This way of putting matters, however, is best accommodated by HEMC; and given the costs to intuition -- and to the general principle of conservatism in theory acceptance -- of spreading the mind out into the world beyond the organism, there seems no reason to reinterpret the situation in keeping with HEC. (2004, p.405)

Similar to the motely-crew problem, this argument appeals to general features of science to argue against extended cognition; or, as Rupert calls it the "hypothesis of extended cognition" (HEC). The idea is that if two theories have same explanatory power, then research should prefer the one that is less radical. Thus, even if extended cognition is a coherent picture for cognitive science, there is still a theoretically more conservative way available to accommodate its insights; what Rupert calls the "hypothesis of embedded cognition" (HEMC). More schematically, the argument looks something like this:

(i) As a general framework for cognitive science, HEMC is a competitor to HEC.

(ii) Conservatism favors HEMC, ceteris paribus.

(iii) Therefore, HEMC should be preferred to HEC.

Rupert's argument requires showing that for any given phenomena the HEMC has the same explanatory power as HEC such that extra-theoretical virtues tip the balance in favor of HEMC.[38] How is it that the HEMC stands as a rival to HEC? Consider how Rupert characterizes HEMC:

> According to the hypothesis of embedded cognition (call it HEMC), cognitive processes depend very heavily, in hitherto unexpected ways, on organismically external props and devices and on the structure of the external environment in which cognition takes place. (2004, p. 393)

Notice the similarity between this statement of the HEMC and the methodological assumption of externalism. That assumption claimed that researchers should turn to internal structures and processes only after investigation was given to how organisms off-load and distribute cognitive processing into worldly structures. The HEMC takes this moral seriously. It looks to study cognitive processes and systems as they depend in unexpected ways on environmental structures. Yet, unlike the HEC, the HEMC does not go so far as to claim that the "epistemological dependence" (causal dependence) leads to environmental structures becoming part of the cognitive

[38] Adams and Aizawa (2008, ch.9) present a similar argument, but it is framed more in terms of inference to best explanation, less in terms of explanatory power.

processes or systems (see Rupert, 2004, p.393). The HEMC denies the metaphysical assumption of externalism.[39]

How do each of the extended views of music cognition fair against this final challenge? As it happens: quite well. This is because each view can argue the ASC to standstill. How is this? Well, note that the ASC requires that there are no dimensions of explanatory power possessed by the HEC that are not also possessed by the HEMC. This is what the ceteris paribus clause of premise (ii) entails. The reason for this is so that the "principle of conservatism" can kick in and tip the theoretical scale in favor of the HEMC.

However, in Chapter 1 it was established that there were at least two other dimensions of explanatory depth that wide explanations have that internalist-friendly narrow explanations can fail to have: theoretical appropriateness and causal depth. Since, as we already saw in Chapter 2, wide computationalism is a species of wide explanations pitched at the information processing level, assuming that the other two forms of extended music cognition also subscribe to wide explanations, which seems a safe assumption, then the ASC fails to deliver its conclusion. It is simply isn't the case that other things are equal, explanatorily speaking. Insofar as each extended music view adopts a wide approach to explanation, as is at least the case for ECMC, then each view can resist premise (ii) of the ASC.

At best, the internalist challenger can argue that there are further

[39] We can see begin to see here the relationship between the internalist challenges. The ASC, in part, draws on considerations from the CCF to draw its boundary between the HEMC and HEC. Such considerations can be put to one side for now, though. For a discussion about the relationship between the internalist challenges see Wilson and Clark (2009) and Menary (2007).

reasons to outweigh theoretical appropriateness and causal depth with the principle of conservatism. But notice that in doing so the argument has shifted. It is no longer about why one should prefer the HEMC to the HEC on grounds of theoretical or explanatory depth. Rather, it is now a more general argument about what makes for good theory selection among the theoretical virtues. Of course, this is an important discussion for the philosophy of science, but it is not one that is directly relevant here. If we have already admitted that one can be a pluralist about explanatory practices in music psychology, there is little further reason to argue that there is some sort of hierarchical ordering of explanatory powers -- a position, I might add, that one should be weary of adopting (see, Dale, Dietrich, and Chemero, 2009).

The more general point is that if the extended views of music cognition can be situated within the wide music psychology, then there is little further reason to try and deflate or collapse the insights of extended views within the internalist framework.

One might respond by saying that the preceding argument requires showing that this is true for each particular view. But this is fortunately not the case. Because the ASC is an a priori argument against extended cognition views, all that needs to be shown is that it fails a priori. There is no need to show that each particular extended view does have the extra explanatory depth of wide explanations when the ASC requires that they can't in principle.

3.3 The Final Scorecard

So, what's the final scorecard? How do the extended views stack up? Well, we saw that Krueger's musically extended mind and Cochrane's extended musical expression accounts commit some version of the CCF, while the ECMC doesn't. It also turned out that Krueger's account and the ECMC had the resources to engage in productive extended musical science, while Cochrane's account came up short. And, we saw that all three views could handle the ASC. So, for those sports fan readers, the scorecard looks something like this: The ECMC has a 2-0-1 record; Krueger's account has a 1-1-1 record; and Cochrane's account has a 0-2-1 record.

The uptake is that, as an approach to cognitive extension within wide music psychology, the ECMC seems to have the most going for it. What should make the ECMC particularly attractive, moreover, is that it achieved its record in spite of granting each of the internalist challenges all the theoretical ground they wanted. The ECMC managed to avoid or accommodate each of the challenges, even while granting them their larger assumptions. Combine these considerations with those of Chapters 2 and the ECMC should seem not only as a viable but desirable approach to extended music cognition.

Bibliography

Adams, F. & Aizawa, K. (2008). The Bounds of Cognition. Malden: Blackwell Publishing.

Adams, F. & Aizawa, K. (2009). "Why the Mind is still in the head". In P. Robbins and M. Aydede's, The Cambridge Handbook of Situated Cognition. Cambridge: Cambridge University Press: 78-79.

Bach, K. (1982). "De Re belief and Methodological Solipsism". In A. Woodfield (ed.), Thought and Object: Essays on Intentionality. Oxford: Oxford University Press, pp.78-91

Baker, M. (2010). "From cognition's location to the epistemology of its nature". Cognitive Systems Research, 11: 357–366.

Balzano, G. (1986). "Music Perception as Detection Of Pitch-Time Constraints". In G. Balzano and V. McCabe (eds.), Event Cognition: An Ecological Perspective. Hillside, New Jersey: Lawrence Erlbaum Associates, Publishers.

Bechtel, W. (2007). "Reducing Psychology while Maintaining its Autonomy via Mechanistic Explanation". In M. Schouton and H. Looren de Jong (eds.), The Matter of the Mind: Philosophical Essays of Psychology, Neuroscience, and Reduction. Malden, MA: Blackwell Publishing.
.
Bernstein, L. (1976). The Unanswered Question. Cambridge, Mass.: Harvard University Press.

Bharucha, J. J. (1999). "Neural nets, temporal composites, and tonality". In D. Deutsch (ed.), The Psychology Of Music. San Diego, CA: Academic Press: 413-440.

Bulter, D. (1989). "Describing the perception of tonality in music: A critique of the tonal hierarchy theory an d a proposal for a theory of intervalic rivalry". Music Perception, 6: 1219-1242.

Burge, T. (1979). "Individualism and the mental". In P. A. French, T. E Uehling & H. K. Wettstein (ed.), Midwest Studies in Philosophy, 4. Minneapolis, University of Minnesota Press.

Burge, T. (1986). "Individualism and Psychology". Philosophical Review, 95: 3-45.

Chang, H. W., & Trehub, S. E. (1977). “Auditory processing of relational information by young infants”. Journal of Experimental Child Psychology, 24: 324-331.

Child, P. (1984). “Review of A Generative Theory of Tonal Music”. Computer Music Journal, 8(4): 56-65.

Chomsky, N. (1957). Syntactic Structures. The Hague: Mouton.

Chomsky, N. (1959). “Review of Skinner's Verbal Behaviour”. Language, 35(1): 26-58.

Chomsky, N. (1965). Aspects of the Theory of Syntax. Cambridge, Massachusetts: MIT Press.

Chomsky, N. (1995). “Language and Nature”. Mind, 104(413): 1-61.

Clark, A. (1989). Microcognition. Cambridge, Mass.: MIT Press.

Clark, A. (2003). Natural-born Cyborgs: Minds, Technologies, and the Future of Human Intelligence. Oxford: Oxford University Press.

Clark, A. (2008). Supersizing the Mind: Embodiment, Action, and Cognitive Extension. Oxford: Oxford University Press.

Cuddy, L. (1991). “Melodic patterns and tonal structure: Converging evidence”. Psychomusicology, 10: 107-126.

Cummins, R. (1983). The Nature of Psychological Explanation. Cambridge, MA: MIT Press.

Cummins, R. (2000). “What are the laws? versus How does it work?”. In F. Keil and R. Wilson, Explanation and Cognition (eds.). Cambridge, MA, MIT: 113-42.

Cochrane, T. (2008). “Expression and Extended Cognition”. The Journal of Aesthetics and Art Criticism, 66(4): 329-340.

Dale, R., Dietrich, E., & Chemero, A. (2009). “Explanatory Pluralism in Cognitive Science”. Cognitive Science: 1–4.

Dawson, M. R. (1998). Understanding Cognitive Science. Malden, MA: Blackwell.

Dawson, M. R., & Yaremchuk, V. (2008). “Artificial Neural Networks that Classify Musical Chords”. International Journal of Cognitive Informatics and Natural Intelligence, 2(3): 22-30.

Dawson, M. (2009). “Computation and Cognition -- and Connectionism”. In D. Dedrick and L. Trick (eds.), Computation, Cognition, and Pylyshyn. Cambridge, MA: MIT: 175-200.

Deutsch, D. (2012). The Psychology of Music. New York: Academic**.**

Devitt, D. (1990). “A Narrow Representational Theory of Mind”. In W. G. Lycan (ed.), Mind and Cognition: A Reader. Cambridge, MA: Harvard University Press.

Dowling, D., & Harwood, J. (1986). Music Cognition. Orlando, Florida: Academic Press.

Fodor, J. (1975). The language of thought. Cambridge, MA: Harvard University Press.

Fodor, J. (1987). Psychosemantics. Cambridge, MA: MIT Press.

Fodor, J., & Pylyshyn, Z. (1981). “How direct is visual perception?: Some reflections on Gibson's ‘ecological approach’”. Cognition, 9(2): 139–196.

Fiske, H. E. (2004). Connectionist Models Of Musical Thinking. Lewiston, N.Y.: E. Mellen Press.

Garfinkel, A. (1981). Forms of Explanation: Rethinking the Question in Social Theory. New Haven, CT: Yale University Press.

Gibson, J.J. (1966). The Senses Considered as Perceptual Systems. Boston: Houghton Mifflin.

Gibson, J.J. (1986). The Ecological Approach to Visual Perception. New York, NY: Psychological Press.

Goldin-Meadow, S. (2003). Hearing Gesture: How Our Hands Help Us Think. Cambridge: Harvard University Press.

Goldstein, B. (1981) “The Ecology of J.J. Gibson's Perception”. Leonardo, 14(3): 191-195.

Griffith, N., & Todd, P. M. (1999). Musical Networks: Parallel Distributed Perception And Performance. Cambridge, Mass.: MIT Press.

Hempel, C. (1965). Aspects of Scientific Explanation. New York, NY: Free Press.

Hammanaka, M., Hirata, K., & Tojo, K. (2006). "Implementing a generative theory of tonal music". Journal of New Music research, 35 (4): 249-277.

Hurley, S. (2010). "Varieties of Externalism". In R. Menary (ed.), The Extended Mind. Aldershot, Hants: Ashgate.

Jackendoff, R., & Lerdahel, F. (2006). "The capacity for music: what is it, and what's special about it?". Cognition, 100: 33-72.

Krueger, J. (2014). "Affordances and the musically extended mind". Frontiers of Psychology, 4(1003): 1-9.

Krumhansl, C. L., & Shepard, R. N. (1979). "Quantification of the hierarchy of tonal functions within a diatonic context". Journal of Experimental Psychology: Human Perception and Performance, 5(4): 579-594.

Krumhansl, C.L, & Kessler, E. J. (1982). "Tracing the dynamic changes in perceived tonal organization in a spatial representation of musical keys". Psychological Review, 89(4): 334-368.

Krumhansl, C.L. (1983). "Perceptual Structures for Tonal Music". Music Perception: An Interdisciplinary Journal, 1(1): 28-62.

Krumhansl, C. L. (1990). Cognitive Foundations Of Musical Pitch. New York: Oxford University Press.

Lerdahl, F., & Jackendoff, R. (1983). A Generative Theory of Tonal Music. Cambridge, MA: MIT.

Longuet-Higgins, H. C. (1979). "The perception of music". Proceedings of the Royal Society, B 205: 307–322.

McCabe, V. (1986). "Event Cognition and the Conditions of Existence". In G. Balzano and V. McCabe (eds.), Event Cognition: An Ecological Perspective. Hillside, New Jersey: Lawrence Erlbaum Associates, Publishers.

McGinn, C. (1982). "The structure of content". In A. Woodfield (ed.), Thought and Object. Oxford: Oxford University Press: 207-58.

Menary, R. (2007). Cognitive Integration: Mind and Cognition Unbounded. Basingstoke, New York: Palgrave Macmillan.

Morton, A. (2002). The importance of being understood: Folk psychology as ethics. London: Routledge.

Piccinini, G, & Scarantino, A. (2011). “Information Processing, Computation, and Cognition”. Journal of Biological Physics, 37 (1): 1-38.

Pylyshyn, Z. (1980). “Computation and cognition: issues in the foundations of cognitive science”. The Behavioural and Brain Sciences, 3: 111-169.

Pylyshyn, Z. (1984). Computation and Cognition. Cambridge, MA: MIT Press.

Pylyshyn, Z. (2009). “Perception, Representation, and the World”. In D. Dedrick and L. Trick (eds.), Computation, Cognition, and Pylyshyn. Cambridge, MA: MIT: 175-200.

Pearce, M. & Rohrmeir, M. (2012). “Music Cognition and the Cognitive Sciences”. Topics in Cognitive Science, 4(4): 468-484.

Radocy, R. & Boyle, D. (2012). Psychological Foundations of Musical Behaviour. Springfield, Illinois: Charles C. Thomas Publisher.

Rowlands, M. (1995). “Against methodological solipsism: The ecological approach”. Philosophical Psychology, 8(1): 1-20.

Rowlands, M. (1999). The Body in Mind: Understanding Cognitive Processes. Cambridge, MA: Cambridge University Press.

Rugg, D., & Curran, T. (2011). “Event-related potentials and recognition memory”. Trends in Cognitive Sciences, 11(6): 251-57.

Rupert, R. (2004). “Challenges to the Hypothesis of Extended Cognition”. Journal of Philosophy, 101: 389-428.

Rupert, R. (2009). Cognitive Systems and the Extended Mind. New York: Oxford University Press

Segal, G. (1989). “Seeing What is Not There”. Philosophical Review, 98: 189-214.

Segal, G. (1997). "Review of R.A. Wilson, “Cartesian Psychology and Physical Minds: Individualism and the Sciences of Mind”. British Journal for the Philosophy of Science, 48(1): 151-156.

Sekuler, R. and Blake, R. (1990). Perception. New York: McGraw- Hill.

Sprevak, M. (2010). "Inference to the hypothesis of extended cognition". Studies in History and Philosophy of Science, 41: 353–362.

Stich, S. (1983). From Folk Psychology To Cognitive Science. Cambridge, MA: MIT Press.

Stufflebeam, R. (1999). "Computation and Representation". In W. Bechtel and G. Graham (eds.), A Companion to Cognitive Science. Malden, Massachusetts: Blackwell Publishers.

Sandra T. E., Bull, D., & Leigh T. A. (1984). "Infants' Perception of Melodies: The Role of Melodic Contour". Child Development, 55(3): 821-830.

Sandra T.E., Bull, D., Leigh T.A., & Morrongiello, B. (1987). "Organizational Processes in Infants' Perception of Auditory Patterns". Child Development, 58(3): 741-749.

Shapiro, L. (2000). "Michael R.W. Dawson, Understanding Cognitive Science". Minds and Machines, 10: 440-444.

Shapiro, L. (2011). Embodied Cognition. London & New York: Routledge.

Shoemaker, S. (1981). "Varieties of Functionalism". Philosophical Topics, 12(1): 93-119.

Temperley, D. (2001). The cognition of basic musical structures. Cambridge, MA: MIT Press.

Thagard, Paul, (2012). "Cognitive Science". In Edward N. Zalta (ed.), The Stanford Encyclopedia of Philosophy: 1-15.

Todd, P. M & Loy, G. D. (1991). Music and Connectionism. Cambridge: MIT Press.

van Gelder, T. (1995). "What might cognition be, if not computation?". Journal of Philosophy, 91: 345-81.

von Eckhardt, B. (1995). What is Cognitive Science?. Cambridge, Mass.: MIT Press.

Ylikoski, P., & Kuorikoski, J. (2010). "Dissecting Explanatory Power". Philosophical Studies 148.2 (2010): 201-19.

West, R., Howell, P., & Cross, I. (1985). "Modelling perceived musical structure". In P. Howell, I. Cross, & R. West (eds.), Musical structure and cognition. London: Academic Press: 21-52.

Wilson, M. (2002). "Six views of embodied cognition". Psychonomic Bulletin and Review, 9 (4): 625-636.

Wilson, R. (1994a). "Causal Depth, Theoretical Appropriateness, and Individualism in Psychology". Philosophy of Science, 61(1): 55-75.

Wilson, R. (1994b). "Wide Computationalism". Mind, 103(4): 351-372.

Wilson, R. (1995). Cartesian Psychology and Physical Minds: Individualism and the Sciences of the Minds. Cambridge: Cambridge University Press.

Wilson, R. (2001). "Two Views of Realization". Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition, 104(1): 1-31.

Wilson, R. (2004). Boundaries of the Mind: The Individual in the Fragile Sciences: Cognition. Cambridge: Cambridge University Press.

Wilson, R. (2010). "Extended Vision". In N. Gangopadhyay, M. Madary, and F. Spicer (eds.), Perception, Action and Consciousness. New York: Oxford University Press: 277-290.

Wilson, R. (2014). "Ten questions concerning extended cognition". Philosophical Psychology, 27 (1): 19-33.

Wilson, R., & Clark, A. (2009). "How to Situate Cognition: Letting Nature Take its Course". In P. Robbins and M. Aydede (eds.), Cambridge Handbook of Situated Cognition. Cambridge: Cambridge University Press: 55-77.

Wixted, J. (2007). "Dual-Process Theory and Signal-Detection Theory of Recognition Memory". Psychological Review, 114(1): 152–176.

Wright, C., & Bechtel, W. (2007). "Mechanisms and Psychological Explanation". In P. Thagard (ed.), Handbook of Philosophy of Science, philosophy of Psychology and Cognitive Science. Elsevier: 32-79.

Images Citations

Figure 1.2. Mozart Sonata K. 331. Image Retrieved from: https://www.google.ca/search?q=Mozart+Sonata+K.+331

Figure 1.3. Krumhansl's Tonal Hierarchy. Image Retrieved from: https://www.google.ca/search?q=krumhansl+tonal+hierarchy