

Vítor Manuel dos Anjos Guerreiro

THE ART OF TONAL SOUND

MUSIC AS A SOCIAL KIND AND NATURAL KIND ACTIVITY

Tese de doutoramento apresentada à
Faculdade de Letras da Universidade
do Porto para obtenção do grau de
Doutor em Filosofia.

Faculdade de Letras da Universidade do Porto

2016

A preparação desta tese beneficiou de uma bolsa de doutoramento atribuída pela Fundação para a Ciência e a Tecnologia (Ref.^a SFRH/BD/71536/2010)

Acknowledgements

I wish to thank, first of all, to Prof. João Alberto Pinto and Prof. Andrew Kania, who both supervised, up to a point, the writing of this work, and who had the patience of following my often volatile and not always clear ideas on the subject at hand. None of the many flaws to be found here are their responsibility.

I am indebted to Desidério Murcho and Aires Almeida for the long-time encouragement, stimulus, and challenge they have provided me.

To Prof. José Meirinhos for all the support, interest and kindness he has always showed me.

To Lucas Miotto for the hours we have spent exchanging opinions and for useful and enlightening commentaries.

I am grateful to Vítor Rua for the many spirited and always fruitful conversations.

To André Pereira and Leonor Coimbra for their patient friendship and for giving me strength in some of my less luminous moments.

To Sergio Mota, for his patience in reading the text and for helpful suggestions.

To Henrique Jesus, for our many fertile and heartening conversations.

To my friend and colleague Susana Cadilha, for making me know that so much of my gloomy spectres were actually more of a trite and shared burden than they seemed.

I am especially grateful to António Poeiras for our long walks and dialogues, during which many of the ideas I present here gained shape.

To Fernando Poeiras, for inspiring me, with a chaplinesque metaphor, to “let the lion come out of its cage”, as well as for his continuous, non-metaphorical support.

To the enduring friendship of Tiago Sousa, as well as his keen interest in both philosophical and musical issues, his curiosity and genuine appetite for discussion. I am thankful for his support and everything I have learned (and keep learning) from our shared experiences.

To Duga, my dog, who (this is intentional and not a grammar lapse) saved my soul from the shadowy shores of a metaphorical but no less desolate Styx.

To Olga, Isabel, Beatriz and all family and friends who gave me relentless support and shared the dubious fortune of putting up with me in better and worse moments.

Finally I want to thank Fundação para a Ciência e Tecnologia for the financial support during the writing of this dissertation.

Resumo

Esta tese procura responder à pergunta “O que é a música?” e, complementarmente, tirar daí implicações para o problema de saber o que é uma obra musical. Essas duas questões são as mais gerais da metafísica da música. Acerca da primeira, trata-se de indagar que condições ou propriedades fazem algo ser música. Acerca da segunda, trata-se de indagar que tipo de entidade as obras musicais (coisas que satisfazem as condições para ser música) são.

À primeira pergunta é dada a seguinte resposta: “A música é arte dos sons tonais”. As objeções ou contraexemplos mais naturais dizem respeito às duas relações de conceitos potencialmente problemáticas: entre música e arte, e entre sons e tons.

A caracterização da música como arte é ambígua entre pelo menos um entendimento funcionalista e um entendimento processualista do que é “arte”. Partindo da distinção entre categorias naturais e sociais, ou melhor, da ideia de que tal distinção, a ser feita, envolve crucialmente a noção de propriedades funcionais, propõe-se uma caracterização bipartida das coisas que são socialmente construídas, ou, em sentido lato, artefactos. Por um lado, temos os itens artefactuais que são adequadamente descritos por uma teoria funcionalista, por outro, os itens artefactuais que são adequadamente descritos por uma teoria processualista.

Aceitando a sugestão do pluralismo conceptual acerca da arte (a ideia de que há mais do que um conceito de arte), propõe-se que entendamos a arte com base num conceito dual: a junção de um conceito funcionalista (estético) metafisicamente mais primitivo, e um conceito processualista (institucional) derivativo. Entender a arte deste modo terá como vantagem quebrar o impasse entre as abordagens funcionalistas e processualistas. Propõe-se que o conceito de música tem uma estrutura igualmente dual que espelha a do conceito de arte.

A abordagem explicativa subjacente à caracterização da música como arte dos sons tonais é de tipo funcionalista: a música é organização intencional de acontecimentos sonoros de modo que estes tenham a capacidade de gerar um tipo peculiar de experiência. A circularidade envolvida em denominar essa experiência como “musical” pode ser desfeita especificando os componentes cruciais dessa experiência: audição acusmática e metáfora estética. O primeiro consiste na experiência dos sons em abstração das causas ou fontes sonoras; o segundo consiste na exploração de um mecanismo perceptivo ou capacidade de escutar os sons em termos de movimento, ação e emoção.

Aqui sugiro que a experiência musical, ou, na expressão de Roger Scruton, experiência dos sons como tons, é um caso especial da experiência estética dos sons.

Não é dada uma resposta positiva à segunda pergunta mas somente uma resposta negativa, no capítulo final: as obras musicais não são identificáveis *tout court* com “estruturas sonoras” abstratas. Esta resposta negativa é dada como uma consequência, para a caracterização das obras musicais, do facto de a música ser uma categoria social/artefactual, essencialmente caracterizada por propriedades funcionais.

Palavras-chave: Música; Arte; Categorias Sociais; Funcionalismo; Institucionalismo.

Abstract

This dissertation seeks to provide an answer to the question: “What is music?” and thence, complementarily, to draw implications for the problem of determining what a musical work is. These are the two most general questions of musical metaphysics. The former is about the conditions or properties in virtue of which something is music. The latter is about the kind of entities musical works (things that satisfy the conditions to be music) are.

To the first question the following answer is given: “Music is the art of tonal sound”. The most natural objections or counterexamples concern the two potentially problematic relations of concepts: that between music and art, and that between sounds and tones.

The characterization of music as an art is ambiguous between at least a functionalist and a proceduralist understanding of what “art” is. Drawing on the distinction between natural and social kinds, or, to be more clear, on the idea that such a distinction, to be made, crucially involves the notion of functional properties, a twofold characterization of socially constructed things, or, in a wide sense, artifacts, is proposed. On the one hand, we have the artifactual items appropriately described by a functionalist theory; on the other hand, the artifactual items appropriately described by a proceduralist theory.

Taking up the suggestion of conceptual pluralism about art (the idea that there is more than one art concept), it is proposed that we understand art on the basis of a dual concept: the conjunction of a functionalist (aesthetic) concept, metaphysically more basic, and a proceduralist (institutional) concept, which is derivative. To understand art in such a way has the advantage of breaking the impasse between functionalist and proceduralist approaches. It is proposed that the concept of music has a similar dual structure, which mirrors that of the art concept.

The explanation which underlies the characterization of music as the art of tonal sounds is of the functionalist kind: music is intentional organization of sound events such that these have the capacity to generate a peculiar kind of experience. The circularity involved in calling that experience “musical” can be undone by specifying the key components of that experience: acousmatic listening and aesthetic metaphor. The first consists in experiencing the sounds in abstraction from their causes or sources; the second consists in the exploration of a perceptual mechanism or ability to hear sounds in terms of movement, action and emotion.

Here I suggest that musical experience, or, in Roger Scruton’s phrase, the experience of sounds as tones is a special case of the aesthetic experience of sounds.

The second question is not given a positive answer but a merely negative one, in the final chapter: musical works are not identifiable *tout court* with abstract “sound structures”. This negative answer is presented as a consequence, to the characterization of musical works, of the fact that music is a social/artifactual kind, essentially characterized by functional properties.

Key words: Music; Art; Social Kinds; Functionalism; Institutionalism

Contents

Acknowledgements	3
Resumo.....	5
Abstract.....	7
Contents.....	9
Introduction	13
1 Natural kinds, social kinds, and artifacts	27
1.1 Kinds of functional properties in music and language	27
1.2 Culturally modulated natural-kind activities	33
1.3 Sortals, mind-dependence and real existence	43
1.4 Disjunction and real definition	51
2 Social ontology and art kind ontology	59
2.1 Function and procedure: causal-agentive art and status-functional art.....	59
2.2 Institutions, games and status-functions	73
2.3 Descriptivism, revisionism and the pragmatic constraint	107
2.4 Should we replace the pragmatic constraint?	121
3 A working notion of the aesthetic.....	137
4 Hearing sounds as tones: toward a natural-kind theory of music.....	159
4.1 The components of musical experience: acousmatic listening.....	163
4.2 The components of musical experience: aesthetic metaphor	183
4.3 Music as the art of tonal sound	215
5 Extant definitions of music: a survey.....	221
5.1 Jerrold Levinson: music as the sonic intensification of experience	221
5.2 Andy Hamilton: music and (non-musical) sound art.....	231
5.3 Andrew Kania: artless music and listening for basic musical features.....	241
6 What musical works are	255
6.1 The Puzzles of Musical Ontology	257
6.2 Musical Works and Social Reality	277
Conclusion.....	319
Bibliography	331

Suppose, however, that I direct my gaze to the far corner of the field, where the long grass is teased by wind, and strange patterns are formed as the light plays against it. It has the look of a witches' Sabbath: I see the dancing forms of the witches, and the black dog, formed of shadow, who bays in their midst. I don't believe that any such event is happening, but that is how I see it, nevertheless.

Roger Scruton¹

Among the various experiences of our senses, tone is the only one that belongs exclusively to life. Light and color, sound, odor, and taste, solidity, fluidity, and gaseousness, rough and smooth, hot and cold, all these are also to be found in non living nature. Only life can produce tones. Living beings out of themselves, add tone to the physical world that confronts them; it is the gift of life to non living nature.

Viktor Zuckerkandl²

¹ Scruton (2012: Ch. 24).

² Zuckerkandl (1956: p. 1).

Introduction

It is not easy to determine the nature of music.

Aristotle³

Definitiveness is as elusive in the metaphysics of music as it is in musical performances.

John Bender⁴

One of my aims in this work is to lay out the components of an aesthetic theory of music. By an aesthetic theory of music I mean a theory that purports to explain what music essentially is, in which some concept (or subconcept) of the aesthetic plays a crucial, ineliminable role. Depending on which concept or concepts are emphasized or deployed for that role, the result will be a specific kind of aesthetic theory. Just as there is no such thing as *the* aesthetic theory of art but a family of theories related by their deployment of aesthetic concepts and the explanatory roles such concepts are expected to perform, the same is true of an aesthetic theory of music. Some aesthetic theories emphasize the concept of a proper function performed by all objects belonging to the class they seek to define, e.g. that of generating *aesthetic properties* that depend on *sounds*,⁵ or on the ability of the object for generating *aesthetic experiences* of sounds in perceivers. As it happens, the theory I wish to defend is perhaps disappointingly unoriginal, since it can be stated using the words of Eduard Hanslick in his work of 1854, *On the Musically Beautiful*: music consists “simply and solely of tones and their artistic combination”⁶ or of “tonally moving forms”⁷. In fewer words, the thesis I’m putting forward is the following:

Music is the *art of tonal sound*.

³ *Politics*, VIII, 5.

⁴ “Music and Metaphysics: Types and Patterns, Performances and Works”, in Chapter 6, *Contemporary Philosophy of Art: Readings in Analytic Aesthetics*, co-authors and eds. J. Bender and H. G. Blocker, Prentice Hall, 1993, pp. 354-365.

⁵ I take this formulation from Zangwill (2007: p. 394).

⁶ Hanslick (1986: p. 28). In the book, this is actually offered as a definition of musical *beauty*. Since he goes on to state that the *content* of music is (nothing but) “tonally moving forms”, the “entire system of tones, with their latent possibilities for melodic, harmonic and rhythmic variety” being the “material out of which the composer creates”, I take it that the first sentence quoted is actually to be understood thus: the (specifically musical) beauty *depends solely* on the artistic combination of tones, i.e., tones and their artistic combination are the sole *dependence base* of musical beauty (which is meant to exclude the non-perceptual relational properties of music).

⁷ *Ibid.* p. 29.

In order to dispel an obvious objection of over-narrowness, it must be cautioned that the ambiguous word “tonal” is not to be taken in the more limited sense of “pertaining to (western) tonality” or to what is sometimes called “the western tonal system” (i.e., major-minor diatonic tonality). Rather, “tonal” is to be understood here in a wider sense, encompassing all organization of sounds into relations of tones.⁸ “Tone” is also not to be understood merely as a sound of regular frequency or definite pitch⁹ (though that is one of the legitimate uses of the word), but as a sound that stands in certain relations to other sounds, in virtue of all these sounds standing in a certain relation to the human *experience* of listening to and organizing sounds. To explain these relations amounts to explain what music *is*. A sound heard as a tone is not merely defined in terms of cycles per second in a regular vibration pattern, any more than the property *being made of bronze* by itself explains why a given piece of bronze is a statue of Napoleon. Sounds that are perceptually indiscernible (token identical) in terms of pitch constitute different tones given certain variations in context; conversely, sounds of different frequencies can count as the same tone, also given certain variations in context – think of “same note”, “same interval” or “same chord” in different tuning systems (e.g. just intonation and equal temperament), and of how standards for tuning instruments have varied within the equal temperament system. In this wider sense of “tone”, being heard as a pitched sound is only one of several components in a sound’s being experienced as a tone or as “belonging to an order of tones”¹⁰ – something for which being pitched is not even a necessary condition, since a purely rhythmic sequence of non-pitched sounds counts, in this sense, as “belonging to an order of tones”. By using “tone” in this manner I’m also drawing on a well-known theory of musical experience developed by Roger Scruton (though Scruton himself does not defend an explicit definition of music) in a series of works.¹¹ I will describe Scruton’s theory in more detail later on, in section 4, and though I will adopt for my purposes some concepts from his theory (particularly, the concepts of *acousmatic experience* and that of *experiencing sounds as tones*), I

⁸ In this wider sense, “tonal” applies equally to music in the major-minor diatonic tonality, modal music, atonal music, 12-tone music, and any music based on “tones”, which includes, of course, non-western music. “Atonal music” is a description of that music that lacks or eschews functional harmony, but even this music is made with tones, i.e., sounds which have rhythmic, melodic or harmonic implications.

⁹ Though it’s easy to confound the two concepts, *pitch* and *frequency* are quite distinct. Differences in frequency which are not perceivable by the human ear are not counted as differences in pitch (an oboist playing a middle A continuously will generate a sound with minute variations in frequency, without it being perceived as a change of pitch). Furthermore, perceivable variations in frequency produced intentionally, as when the same A note is continuously played on a bowed instrument using vibrato, will nevertheless count as episodes in the playing of a single note. Though all sounds have frequency, not all sounds are “pitched”, such as the sounds of untuned percussion instruments or everyday noises (sounds with irregular vibration patterns). Frequency is an intrinsic property of sounds; pitch is a relational property.

¹⁰ The other components in this experience of implicit organization being *melody*, *harmony* and *tone color*.

¹¹ See in particular, Scruton (1983; 1997; 2009).

don't wish to give the impression that I endorse every aspect of it. (I'll have something to say about Scruton's anti-realism, as it reflects on his discussion of the nature of music, though I wish to leave the door open for anti-realist interpretations of the above definition.) We'll see this in more detail when I explain the theory and some of its more contentious aspects.

The view of music I'll expound here is not what is sometimes called a *formalist* view, which is commonly associated with Hanslick. (Whether and to what extent Hanslick himself had a formalist view of music or whether he consistently held such a view, is a matter of dispute, into which I won't delve here¹².) It is not always clear what is meant with "formalism", and this is due to the absence of a clear-cut distinction between those properties of an artwork that are "formal" and those that are not. The view I favor is not "formalist" to the extent that I endorse contextualism about aesthetic qualities; that is, in my view, the aesthetic qualities of an object depend as much on the relations it has with other things as it depends on its physical structure and perceptual properties taken in abstraction from its surroundings. Furthermore, I believe we can't think of our aesthetic judgments in a given situation as disconnected from our aesthetic judgments in other situations. Our talk of aesthetic judgments sometimes gives the impression that whenever we judge something aesthetically, we "begin from scratch", so to speak. Also, to see something as endowed with a certain aesthetic quality may require us to think of that thing in terms of *something else*, so that if we completely abstract the object from its relations to other things, and our experience of the object from the rest of our experience, we won't be able to discern aesthetic qualities in it. If this is true, then the thesis that music is essentially aesthetic is compatible with the view that to hear sounds as music we must hear them in terms of something else. The thesis that music is essentially aesthetic, as I understand it, says that hearing sounds as music is *one* peculiar mode of hearing sounds *aesthetically*. Non-musical cases of "aesthetic listening" are the cases where one appreciates the aesthetic qualities of words in spoken poetry and prose; the appreciation of "soundscapes" that may be a part of say, garden or fountain design; as well as the case of *non-musical sound art*, which I'll approach when surveying Andy Hamilton's account of music in the final section. An aesthetic thesis in the lines just sketched is not anti-realist, that is, it is not the thesis that for something to be music is for it to be *experienced* as music. My approach is production-centered: that a certain event (or type of event) is music is both a function of its phenomenal properties as well as of the

¹² I'm thinking about the contrast between what Hanslick explicitly says when theorizing about music in the "philosophical mode", and while writing as a music critic, where he makes abundant use of predicates for expressive properties that he explicitly denies are essential to music in his more theoretical work.

intentions of those who *organize* it (composers and performers), that is, who intentionally endow it with the relevant properties. I discuss these topics with more detail in sections 4 and 5.

A certain tradition of reiterating the idea that it does not really matter whether a certain piece is music or not has become more or less conventional wisdom.¹³ Take the case of Cage's *4'33''*. Many people agree that it is an important work of art, one that invites us to think about the role of silence in music but also invites us to notice the aesthetic qualities of those sounds usually considered "ambient" and wholly external to music (the piece is not really a *silent* piece, since its performances are supposed to be constituted by the unpredictable ambient noises that arise during a performance of it). Why should it matter whether it is music or not? And why should it matter whether a piece for twelve radios is music? This is one factor that contributes to a relative lack of interest in a definition of music, and accounts for the relative scarcity of attempts to define music in the literature, when compared with the amount of work done on the definition of art. Take, for instance, the following remark by Phillip Ball in his book *The Music Instinct*:

"Whatever you think of the twentieth-century Western avant-garde, it served as a reminder that this sort of exercise in definition is futile. Music can be made from mistuned radio sets, from the incidental noises of a concert hall, from the hum of machinery. No one says you have to like it."¹⁴

But this is surely too hasty. That something is presented as music in an artworld context, and came to be accepted as such may count as (defeasible) evidence in favor of its being music, but it doesn't *show* that it is. That we may *like* it or find it aesthetically valuable and even consider it to be an interesting or important case of *art* does not settle whether or not it is music. And think about the idea that it is not really important whether something is music. Granted, it may be the case that the concept MUSIC is irremediably arbitrary, but that the friend of the avant-garde tends to make a big deal of the *denial* that pieces such as *4'33''* or *Imaginary Landscape n°4* are music is not consistent with the idea that it doesn't really matter whether we count them as music or not. In other words, people seem to think they are saying something interesting or relevant when they say that a certain artwork is a piece of music. But if "music" is an empty concept, nothing interesting is really being said about, say, *4'33''* by saying that it is music. Even granting that any "organized sound event" can be music, this does not show by itself that sound *simpliciter* (or silence *simpliciter*, for that

¹³ See, for instance, Julian Dodd's discussion of *4'33''* in a lecture at Tedx: <https://www.youtube.com/watch?v=WTCVnKROlos>

¹⁴ Ball (2010: p. 10).

matter) is music, in which case the concept of music would be empty or inoperative. Maybe the fact that a tradition of music making can include virtually any kind of sound event (which includes silent pieces) or even purely conceptual works that are not to be “sounded” can be explained by some relevant factor pertaining to the nature of music or the nature of art practices. By dismissing the question from the outset as “not really important” we deprive ourselves of that kind of explanation.

One important aspect of a definition concerns its motivational adequacy. A definition is motivated when “there is some background theory or description of a well-established practice which explains why it is reasonable for one to be interested in whether a thing meets the proposed condition or why there is a recognized category of things which satisfy that condition.”¹⁵

In other words, a definition is motivationally adequate when we have good reasons to be interested in things of the relevant kind. Now, as Stephen Davies points out: “Music is ancient, pan-cultural, and, given the spontaneous emergence of song in children, virtually universal.”¹⁶ Music is also multivalent – it can serve an indefinite variety of functions – and it has (both philosophically and empirically) puzzling features (one of which that has perhaps received most attention is the way we vividly experience music as having emotional qualities; the fact that a mindless, non-sentient sequence of sounds with no semantic meaning can be experienced as sad, somber, joyous, melancholic, etc.)¹⁷ and having these sorts of puzzling features motivates the search for a good explanatory theory about the *kind* of thing that musical entities are and what makes them *musical* instead of something else (regardless of those properties being referred to with the adjective “musical” instead of any other noise or squiggle).

This brings us to another perplexing feature that music shares with art kinds¹⁸ in general: to what extent is the nature of art kinds independent of how we use concepts and words for art kinds? Do we refer to art kinds in the same way we refer to natural kinds, where an originating use of a term for a natural kind is introduced by *ostension* – “let’s call things of *this* kind ‘copper’” – and reference is secured even if we completely ignore the essential individuating properties of the natural kind in question? (Copper is copper, and it has the properties it has, and we are successful in referring to it, no matter what we call it, and no matter what mythological stories one comes up with

¹⁵ J. Kingsbury & J. McKeown-Green (2009: p. 2).

¹⁶ Every culture has something that “westerners” tend to recognize as “music”, even if there is no equivalent in that culture to what is known as the “western” concept of music or a “term to encompass music as a total phenomenon”.

¹⁷ Ball, *op. cit.*, p. 2: “It is deeply puzzling why these complex mixtures of acoustic frequencies and amplitudes make any sense to us, let alone why they move us to joy and tears.”

¹⁸ The notion of *art kind* here is meant to apply both to things falling under non-sortal concepts expressed by mass or non-count nouns, such as *art* and *music*, and things falling under sortal concepts, expressed by count nouns or expressions, such as *artwork* and *musical work*, but also *painting* and *string quartet*.

to explain its existence.) But are the properties of art kinds *independent* of whatever information is associated in our heads with art kind *terms*? Can we make genuine *discoveries* about art kinds? This question arises whether one endorses the view that art kinds can be given definitions or whether one is an anti-definitionalist about art kinds.

One aspect concerning the *reference* to art kinds is sometimes presented as a reason to doubt that a definition of music is feasible, desirable or useful (though, oddly, the same doubt does not arise for *depiction* or *painting*), namely, that

“Most languages of the world don’t have a term to encompass music as a total phenomenon. Instead, they often have words for individual musical activities or artifacts such as singing, playing, song, religious song, secular song, dance, and many more obscure categories.”¹⁹

“When I refer to ‘music’ I mean something that we, as members of a modern Western society, would recognize as music. But we must immediately appreciate that the concept of music is culturally variable; some languages lack words or phrases to encompass music as the total phenomenon that we understand in the West. Instead, they may have words for specific types of musical activities, such as religious song, secular song, dance, and the playing of particular instruments.”²⁰

Whatever the strength of these remarks, I don’t see them as counting against the motivation for a definition of music but as counting for that motivation. Despite the fact that where different cultures have terms to refer to music, their use of those terms may differ substantially from the way we use the term “music”, it is still true that: 1) there is something which we primarily call “music” in our own cultural tradition, and 2) all cultures have something that we feel compelled to call “music” even if there is no equivalent term in the target culture. That the use of the music concept is culturally variable is compatible with there being one interestingly unified kind of sound-oriented activity which we pick out with that concept (by using it in the “western” sense of a “total phenomenon”, by contrast, say, with “locally inflected” uses of the concept – Nettl, for instance, speaks of the way the word “muziki” was used in Persian culture, to denote instrumental playing but excluded such things as the singing of sacred texts). In fact, such value-laden uses of the concept will be more or less unavoidable, since “local conceptions” of the musical phenomenon will emphasize those properties of music that are most salient and valued given the specific context, and as we have seen, the most salient or valued properties will likely not coincide with the essential

¹⁹ B. Nettl (2005: p.21).

²⁰ S. Mithen (2005: p. 11).

properties. The properties which are most salient and valued in religious depictions will differ across cultural contexts; but none of this will have to do with the properties in virtue of which a depictive representation is a depictive representation, no matter how “culturally inflected” depiction is.

The use of concepts is culturally variable. Music, depiction, art, are all culturally variable. But they are also distinct from the social practice of using concepts to capture and enshrine specific properties of music, depiction and art. So, even if Nettl is right that “there does not appear to be an interculturally valid definition of music”²¹ (that is, a use of the music concept which is accepted by all cultures), this doesn’t show that “music” can’t be used (as “depiction”) to pick out a non-arbitrary, cross-cultural activity, which is relevantly unified and partly explains the diversity of uses of the concept, by establishing a *primary* use in light of which all *extended uses* are to be understood. But Nettl seems to be wrong about the “intercultural validity” issue, for even if the “wide” concept of music as a “total phenomenon” is specific to western culture, that concept has come to be accepted by other cultures as well. Analogously, if the “universality” of the concept of *depiction* was never in doubt just because convergence in the use of *some* concept for that kind of thing took its time to arise, then it is puzzling that the universality of the music concept should be doubtful, and nevertheless it raises doubts.

What is so special about the concept of music, analogously to what happens with the concept of art and unlike the concepts of depiction or painting, that it makes us unsure whether presence of the concept is a necessary condition to the presence of the phenomenon? Such questions suggest a relevant similarity between the concepts of art and music which is absent between, e.g., the concept of art and that of painting, namely, the concept of music seemingly carries a heavier metaphysical luggage than at least some concepts for individual art forms. In light of this fact, considering extant theories of art may help to shed light on a plausible route to explain the nature of music, so that we may learn about the latter by seeing whether a particular theory about the former applies to it or not, e.g., whether or not the fact that something is music may be explained institutionally or in terms of another kind of theoretical approach.

This work is divided into six chapters. In the first I explore the notion of functional properties and the different kinds of such properties that determine two major and overlapping categories of artifactual items: causal-agentive artifacts and status-functional artifacts. This division corresponds to the distinction between non-institutional socially constructed things and institutional socially

²¹ B. Nettl (2005: p. 45).

constructed things. The language of “causal-agentive” and “status-functional” is borrowed from John Searle’s (1995) taxonomy of functional properties: functions are either agentive or non-agentive, i.e., either imposed by agents or “discovered” (e.g. the biological functions of bodily organs). Within the class of agentive functions, causal-agentive functions are those functions an item may perform in virtue of its physical structure or causal powers (these are the functions of typical artifacts such as screwdrivers and corkscrews). Status-functions, on the other hand, are functions an item performs in virtue of a collectively recognized status. The interest in artifactuality is prompted by the notion that both art and music belong to the category of human-made functional things, whatever else they may be, and that this fact is crucially relevant to the ontology of art and music kinds (such kinds are those things falling under the concepts of art and music and which we refer to using sortal concepts, like “artwork”, “musical work”, “painting”, “string quartet”, and so forth). In line with the analogy between the concepts of music and art, I survey different definitions of art and explore ways of classifying definitions of art, such as Stephen Davies’ distinction between functionalist and proceduralist definitions and George Dickie’s scheme of “natural kind theories of art” (NKTA) and “cultural kind theories of art” (CKTA). In this work, I focus particularly on functionalist and proceduralist theories of art, though Davies’ classification also includes the category of historical theories, since my primary aim is to draw important consequences for the definition of music, and no philosopher of art that I know of has proposed an historical definition of music. Notably, the chief proponent of an historical definition of art, Jerrold Levinson, argues for a functionalist definition of music.

My purpose is to draw a correspondence, on the one hand, between those things susceptible of a procedural definition and status-functional artifacts, and, on the other hand, between those things susceptible of a functionalist definition and causal-agentive artifacts. To the extent that art and music depend on “natural-kind activities” (activities that are possible or conceivable in the absence of conventions even if they are always “culturally modulated”) they will be defined in functionalist (causal-agentive) terms. For instance, aesthetic functionalism says that artworks are artifacts with at least an aesthetic function. Something has an aesthetic function if it is meant to exhibit aesthetic properties or to be experienced aesthetically.²² Aesthetic functions are thus a case of causal-agentive function. Also, to experience things aesthetically is a case of natural-kind activity (it can be inflected by conventions or “culturally modulated” but its existence is not dependent on

²² An example of what an aesthetic function is is given by Zangwill (2007: p. 99): “For example, a painting of Mondrian’s mature period has the aesthetic function of having the aesthetic properties of elegance and boldness in virtue of a certain arrangement of black, white and primary colored rectilinear shapes.”

conventions, the way, e.g., *games* exist in virtue of constitutive conventions). However, the concepts of *artwork* and *musical work*, to the extent that musical works are kinds of artwork, also have an element of status-functionality. So in this sense, both art and music are simultaneously causal-agentive and status-functional, though status-functional forms of either can only exist given the existence of causal-agentive forms.

An analogy is drawn between “perception theories” of depiction, which conceive of the power of pictures to represent as relying crucially on a “natural-kind activity”, namely, our ability of “seeing-in” (in other words: depictive representations exploit our ability to notice resemblances between our experience of configurations of lines and pigments and our experience of the represented items). In a like manner, the ability to experience things aesthetically would be the natural-kind activity on which art qua described by the causal-agentive concept depends; and the ability to *hear things in sounds* (namely, to hear motion, action and emotion *in* sounds) would be the natural-kind activity on which music qua described by a causal-agentive concept depends. Both these “natural-kind activities” are described in greater detail in chapters three and four.

In the second chapter I explore the methodological thesis of the Pragmatic Constraint on the metaphysics of art kinds. According to the Pragmatic Constraint, metaphysical theorizing about art kinds must be constrained by actual artistic practice, i.e., by “rational reflection” (to use D. Davies’ (2004) phrase) on the beliefs of “practitioners” (namely, artists and art critics). The Pragmatic Constraint marks a sort of divide between two philosophical approaches to the metaphysics of art: revisionism and descriptivism. Revisionists are persuaded that the metaphysics of art kinds is fundamentally no different from the metaphysics of everything else, so that there is no special methodological approach one should adopt towards art kinds. The descriptivist, however, argues that our metaphysical picture of art kinds cannot run too much adrift from how such things are conceived within the relevant practices, so that the pre-theoretical beliefs and intuitions of “practitioners” must constrain metaphysical enquiry. The central idea in descriptivism is that the dependence of art kinds relative to social agents gives the latter epistemic privilege on the nature of such items, such that there are no fundamental *discoveries* one can make about them that would prove our “pre-theoretical picture” of them utterly wrong.

The idea that there are no discoveries, whether empirical or *a priori*, to be made about social kinds seems plainly wrong: questions about the origins of language, art and music are empirical questions to which we have no ready-made answers; and questions such as exactly what the difference is between pictorial and semantic representation, what mechanisms are at work in both cases, are largely *a priori* questions to which we have no ready-made answers either, despite the fact that the existence of social kinds depends on the actions and directed states of social agents.

Add to this the fact that upholders of the Pragmatic Constraint, as much as “revisionists”, end up with contrasting theories of the nature of art kinds and each of them seems to deny one or other deeply held belief or “folk intuition” about their subject matter (e.g. that artworks are action-tokens rather than the products of such actions), and one wonders whether the Pragmatic Constraint really makes any difference. Different theories will claim to be more “in tune” with pre-theoretical beliefs and intuitions by emphasizing one or other “intuition” over another. For instance, Platonists who conceive musical works as eternal, non-creatable abstract objects will focus on the issue of “repeatability” over other equally “intuitive” characteristics, while others will privilege different and incompatible “intuitions”. But this happens precisely because our “intuitions” exhibit conflicts and inconsistencies. This prompts the question: if our pre-theoretical beliefs about art kinds have inconsistencies, how can they constrain art kind metaphysics? What appears to inevitably happen is that different theories will emphasize different “intuitions”, and each rival theory will accord with the Pragmatic Constraint at some point, and depart from it at another.

In chapter three I give an account of what it is to experience things aesthetically. Such account is meant to be neutral as to the question of the reality of aesthetic properties (properties or qualities such as *beauty*, *ugliness*, *elegance*, *garishness*, *delicateness*, *unity*, *coherence*, *sublimity*, *vibrancy*, *etc.*) and the truth of theories concerning their nature. The proposal (or “working notion”), largely based on Eddy Zemach’s contribution to the subject, says that we experience the world aesthetically to the extent that we experience its phenomenal properties as *desire-mediated*, that is, we interpret phenomenal properties in our environments as impacting our sensibilities and interests. So, for instance, a configuration of lines and colors will be “coherent”, rather than “noisy”, to the extent that its diversity is not such as to overwhelm our cognitive capacities and relative to our desire of understanding the world around us, our desire that the world is intelligible to us. The same goes for our descriptions of objects and events (either natural or artifactual) as *somber*, *gloomy*, *dynamic*, *vibrant*, *delicate*, *sublime*, *elegant*, *grotesque*, *monstrous*, *disgusting*, *ludicrous*, *sinister*, *awesome*, *etc.* This proposal contrasts with views of the aesthetic according to which there is a separate kind of experience, the aesthetic experience, which is characterized by some distinctive feature, namely, hedonic, epistemic or axiological, severing that experience from the regular flow of our everyday experiences. According to the proposal, *all* of our experiences have an aesthetic component, in virtue of the fact that we are the kind of beings capable of forming beliefs and desires. As limited, finite, desire-endowed creatures that are “always in the process of losing everything we have” (Sartwell: 2004, p. 4), we are able to discern an aesthetic surface in the appearance of things. This underlies our ability to exploit *cross-domain mappings*, that is, to enact in perceptual experience a “transfer of concepts derived from one sensory modality to experiences derived from another

sensory modality” (Young: 2014, p. 19), as when, e.g., we describe sequences of sounds as “delicate” or “brittle”. In other words, our ability to “see things” *in* configurations of lines and colors and to “hear things” *in* organizations or combinations of sounds, to experience those things *in terms of something else*, is a crucial component of our experiencing the world aesthetically, i.e., as meaningful in the sense of affecting our desires (only as such can a landscape be described as, e.g., *desolate* or *overbearing*). Art and music emerge when we intentionally and purposefully exploit this ability of ours in the making of objects which have at least as one of their functions to trigger or exploit this ability.

In chapter four I discuss Scruton’s theory of musical experience in its two main components: *acousmatic listening* and *aesthetic metaphor*. According to the acousmatic thesis, it is a condition of experiencing sounds as music that they are listened to in abstraction of their causal sources. This is not to say that in hearing sounds as music we must be unaware of their origin in physical sound sources, such as musical instruments. The thesis is rather that in experiencing sounds as music we experience them as exhibiting an organization, an “inner life”, which is independent of our awareness of how sounds are physically produced. This is partly a consequence of our ability to exploit cross-domain mappings in perceptual experience. I consider Stephen Davies’ objections to “metaphoricism”, the thesis that descriptions of music in terms of height, motion and emotion are all metaphorical descriptions. Davies endorses a resemblance theory of musical expressiveness and as such considers attributions of emotional expressive properties to music as literal, not metaphorical. In his view, attributions of emotion qualities to music expand the literal meaning of the predicates but do not constitute a metaphorical use of such predicates. This discussion is ultimately about the issue of whether or not attributions of emotion qualities to music consist in the attribution of the *same* property denoted in uncontroversial applications of emotion predicates. According to metaphoricism, descriptions such as “the music is sad” are metaphorical devices that either attribute no property at all (Scruton) or gesture toward an aesthetic property of music which for some unclear reason we feel compelled to describe as “sadness” (Zangwill). I argue that whether one adopts a metaphorical or a “literal expansionist” view of descriptions of musical experience, the underlying mechanism to both is still our ability to exploit cross-domain mappings between the domain of sounds and the domain of spatial movement, action and emotion (whether the experience of music *resembles* experience of spatial movement, action and emotion in some way, or music has properties which for some psychological reason we feel compelled to describe with movement, action and emotion predicates). What this means is that musical experience is *musical* partly in virtue of our ability to experience sounds *in terms of something else*: when we hear movement,

action or emotion *in* sounds, whether or not the predicates of movement, action and emotion are being used metaphorically to describe the experience.

In chapter five I survey three approaches in the literature to the definition of music, namely, those of Jerrold Levinson, Andy Hamilton and Andrew Kania, underlining their strengths and limitations. Ultimately, I hold that the most promising attempt at definition is Kania's disjunctive strategy, which is in line with my twofold characterization of the "music concept" as the combination (overlapping) of a fundamental causal-agentive concept and a status-functional concept of music, mirroring the same twofold structure in the "art concept". According to Kania, music consists of sounds intentionally produced to either exhibit "basic musical features" (in other words, rhythmic, melodic and harmonic organization) or to be listened for such qualities even if these are absent (i.e. to be listened for against a background expectation of finding tones and tonal organization). The first disjunct of this definition is very close to my initially proposed definition of music as "the art of tonal sound", except for the fact that Kania rejects the idea that everything which is music is also art, the same way that not all depictions are artistic depictions. One purported example of this (something which is music but not art) is that of someone practicing scales and arpeggios. Despite this, I argue that the conceptual connections between music and art are much deeper than those between depiction and art, so that it is not inappropriate to define music as the art of tonal sound and to see the musical status of purported examples of non-art music as derivative. In other words, we conceive of music's ingredients, such as scales and chords, as musical entities *because* we have the causal-agentive concept of (art) music.

Finally, in chapter six, after explaining the ontological perplexities raised by musical works, I draw some constraints imposed on musical ontology (the problem of what kind of thing a musical work is and what makes a sound event *S* an instance of musical work *W*) by a definition of music in aesthetic-functionalist terms. I argue that in light of a functionalist explanation of music's nature (to the extent that functional properties essentially involve contingent existents, namely, the intentional states of social agents), musical works can't be identified with sound structures *tout court*. Sound structures are acoustic patterns, that is, patterns of vibrating air molecules or pressure waves. Musical works are made of things such as tonal relationships, things which are not to be identified with any acoustic patterns, no more than a description of the chemical composition of a painting explains why it is a depiction of such and such. Musical works belong in *styles* and presuppose networks of relations in the absence of which the "sound structure" *per se* would not even be intelligible as a musical work, in a similar way that only in the context of a language can an acoustic pattern count as the uttering of a *word*. I take up the notion of *indication* as it figures in Jerrold Levinson's theory of "indicated structures" / "initiated types" and Robert Howell's further

elaboration of it, and suggest to supplement that notion with Searle's account of the creation of institutional entities: to create an artwork is to make it the case that one or more tokens of a "structure" (X) count as bearers of the artwork status-function (Y) in the appropriate context (C). The purpose is to provide an account of how musical works are creatable, while remaining agnostic as to whether there are such things as eternal or creatable abstracta. The creation of a musical work is much like the creation of an institutional entity, even though there is a deep asymmetry between fundamentally causal-agentive phenomena such as art and music, and status-functional artifacts (e.g., words, languages, games). To create an institutional entity is not a matter of adding to the world's ontology. When a painter creates a painting, he merely reconfigures the disposition of pre-existent materials such as pigments. Nevertheless, the world is changed and something new is available to us when he does so. When an individual is ordained a priest, no numerically distinct individual was introduced in the world, yet the world is changed in some respect. And while creation certainly involves change, it is not clear why it should involve the introduction of numerically distinct entities in the world. Analogously, when a composer creates a musical work *W* at time *t*, the world before *t* is such that it is an open possibility that a sound event instantiates the same "sound structure" as a sound event after *t* which also counts as a performance of *W*. But this possibility does not make it the case that *W* already existed before *t*, the same way that the Church did not pre-exist the social actions that established it, just because its existence was an open possibility before that time or because no new items were added to the world's fundamental ontology. To create something is to make it *actual*, and not making it the case that its existence is an open possibility in the world.

1 Natural kinds, social kinds, and artifacts

1.1 Kinds of functional properties in music and language

To be a sound of definite pitch is not a sufficient (nor necessary) condition for being a musical sound²³ (or a “tone” as explained in notes 8 and 9 of the Introduction) any more than exemplifying the acoustic pattern of the word “sound” is a sufficient condition for a sound event to count either as an utterance of that particular word or as a word at all (the point being that independently of certain human practices, no sound is a word, whatever acoustic pattern it exemplifies). By making this analogy, though, I’m not implying that some sounds become music *the same way* other sounds become words, that is, I’m not implying that the properties in virtue of which a given sound event counts as a stretch of music are the *same* properties in virtue of which a given sound event counts as an utterance of a word. Music is *not* language; it is not even a language *of sorts*²⁴ (except in a metaphorical use of “language”)²⁵. This point can hardly be over-emphasized. The analogy between the musical example and the linguistic one is that in each case we can imagine a pair of (perceptually) indiscernible counterparts, one of which fails to stand in the appropriate relation to a human practice and so does not qualify, in the one case, as a stretch of music, and, in the other, as a word. In both cases the key properties that distinguish stretches of music and word-tokens from their imagined acoustical counterparts are relational, not inherent properties of the sounds; nevertheless, there is a considerable difference between the musical case and the linguistic one. In fact, it is precisely that difference, as I see it, that ultimately distinguishes how an aesthetic theory and an institutional theory account for the difference between a stretch of music (or an artwork) and its non-musical (non-art) indiscernible counterpart. We can summarize this difference in the contrast between convention and perception. “The meaning of a piece of music is given by

²³ The expression “musical sound” is ambiguous between “a music-like sound” (something that *sounds like* music but might not be music, e.g. birdsong) and “a sound which is music”. I here mean the latter.

²⁴ Danto (1973: p. 15): “The space between art and reality is like the space between language and reality partly because art is a language of sorts.”

²⁵ If the expression “musical language” is an apt metaphor, this is in virtue of the fact that music is *language-like* in a few limited ways: it is made by recombining the same discrete, repeatable elements, to form an endless variety of other repeatable items. Yet there is no mapping from the elements of language to the elements of music. In music there are neither words nor sentences. While music can acquire conventionally arbitrary meanings (think of national anthems or themes representing characters in program music), this is something that is superadded to the music: stretches of music representing characters in a narrative are already *music* independently of their imposed representational function, but this differs radically from the linguistic case. Words are not words independently of being (used as) contingent vehicles of meaning, neither are the music-like properties of words (their aesthetic properties) “imposed”. At this point, the analogy between music and language breaks down.

perception, not by convention”, Scruton has claimed.²⁶ Actually, matters are a bit more complicated, since perception may be informed by convention, so that meaning can be given both by perception and convention. For instance, words have conventional meanings, but my experience of a sound as a word, as opposed to meaningless noise, is also perceptual, in the sense that the phenomenology of both experiences (of a sound as a word and of a sound as meaningless noise) differs. In hearing a sound as a word, I don’t first become aware of a noise to which I then consciously apply a convention in order to generate a word. Instead, I simply perceive the sound as a word. In this aspect, words are not different from music, in that sounds perceived as music don’t *become* music after we consciously apply a convention to the sounds; we simply perceive certain sounds as music, spontaneously differentiating them from mere noise. The analogy can be stretched further: neither the linguistic case nor the musical case require command of a *concept* for words or music. Human beings begin to perceive certain sounds as words *before* they come to grasp the concept WORD. Likewise, to perceive a sound sequence as a *musical* sequence the concept MUSIC is not required. Human beings exhibit musical behaviors (e.g., singing and dancing) before they come to master the subtleties of any concept for music. To perceive sounds as music is no more a matter of having one’s perception guided by the concept MUSIC than to perceive sounds as words is a matter of having one’s perception guided by the concept WORD. It is not a matter of consciously or subconsciously applying those concepts to the sounds one hears.

The crucial difference between the linguistic case and the musical case is that conventional meaning is an essential component of what it is to be a word whereas any conventional meaning that a stretch of music may acquire is not constitutive of its character as music. Scheherazade’s theme, from the same named Symphonic Suite, composed by Rimsky-Korsakov in 1888, which consists of a melody in the violin with harp accompaniment, would still be music if it had been created as part of a piece with no “program”, no connection to the folk tales of *One Thousand and One Nights* or any other narrative, with no intentions of creating a “musical illustration”. The main theme of Beethoven’s *Fifth Symphony* in C minor Op. 67 would still be music if it had never become associated in people’s imaginations with “fate knocking at the door”. But these are the obvious examples. A slightly more interesting case is that of a movement in sonata form, where a melody which constitutes the minor key statement of the second subject in the recapitulation is correctly heard as a restatement of a melody which occurred in the major mode in the exposition.²⁷ Suppose this is a movement in a work with no title or program. No musical “illustration” of an

²⁶ Scruton (1997: p. 171).

²⁷ I recalled this example from Davies (2012a).

extra-musical narrative is intended; no imaginative associations with extra-musical events are brought into play. Yet, the “meaning” of that melody *as* a component of minor key sonata form is conventional: nothing counts as a movement in sonata form independently of what the conventions for that musical form are. However, the status of that particular melody as a musical entity is independent of its belonging to a work that fits the conventions for minor key sonata form. The same melody could occur in a completely different context, where no conventions for sonata form were in place. Of course, there are many other conventions involved in the sonata form example; for instance, a tuning system allowing for modulation between keys, on which the conventions for sonata form depend. (There are other conventions at play here, such as major/minor diatonic tonality.) Yet, music comes *before* all these conventions, pretty much the same way that our ability to perceive certain sounds as words comes before the conventions about, say, the structure of sonnets. The existence of music is conceivable in the absence of conventions about musical form or the tuning of instruments (which is not to say that in such circumstances music has no form whatsoever and instruments are not tuned). The existence of words, on the other hand, is not conceivable without convention, that is, without convergence in social behavior (this is perhaps a pleonasm, since in being *social*, such behavior is already characterized by convergence).²⁸ For a word to exist *just is* for a sound (or inscription) to be used by a group of people as a contingent vehicle of meaning,²⁹ such that it can be combined with other sounds (or inscriptions) to generate sentences, which are also contingent vehicles of meaning.

Many musical entities owe their existence to conventions, or to a network of interlocking conventions, as is the case with movements in sonata form. But the musical nature of such entities is not itself the result of a convergence in behavior. Take a simple example: equal temperament is a conventional device that, among other features, allows for free modulation between the different keys. But the fact that it allows this is not itself the result of convergence in behavior or convention. Rather, it is the musical effect that explains the convention: people converged on the use of equal temperament *because* it delivers the desired result of free modulation between keys. *Musical experimentation* is what ultimately guides convergence in behavior (though outside forces may be equally decisive: take, e.g., counter-reformation and polyphonic singing). The inverse happens in the linguistic case: “green” would do just as well as “red” to refer redness, as long as everyone else

²⁸ Here it must be noted that while all conventions are convergences in behavior, not all convergences in behavior are conventions.

²⁹ “Contingent” in the sense that collective behavior could have converged on any other sound (or inscription) to perform the same role. That “red” is the English word for red is a contingent fact; things could have been such that “green” was the word for red and “red” the word for green.

converged on that use. Here convergence of behavior explains why “red” refers to redness. People did not seek to converge on the use of “red” because of its pre-conventional efficiency in referring to redness. Instead, “red” means red *because* there is a convergence in its use to refer to redness. Sounds and inscriptions become words by becoming *tools* in a human practice, i.e., in coordinated human behavior. They acquire artifactual functions that are a consequence of convergence in behavior. Words have no existence, no reality independently of such context.

Conventions need not be the result of explicit decisions. Though sometimes a word is introduced in a language by stipulation, the majority of words we use are not the result of stipulation but of unintended (unplanned) convergence in behavior, given the advantages of convergence in representing the world. Words are artifacts. Artifacts are defined partly by their functions, which in many cases (though not all) derive from intentions. Where the function itself cannot be performed apart from the convergence in behavior, we can say that the convention is *regulative*, as opposed to *constitutive*. Two examples where a convention is merely regulative. a) In a context where everyone else uses the system of just intonation, we can conceive of an isolated musician writing pieces for a system of equal temperament (this is actually how equal temperament came to be adopted: it began by being advocated by isolated theorists such as Vincenzo Galilei, until it became the predominant tuning system); widespread adoption of equal temperament merely *regulates* (one could appropriately use the metaphor *modulates*) an activity that is conceivable before the convention is in place. b) In a context where there are no rules concerning the side of the road one drives on, there can still be driving or road traffic (however dangerous); in such a context it is conceivable that some people begin to drive consistently on the right (or the left). Convergence in behavior generates a convention, but the convergence is not constitutive of the behavior thus regulated (there can be right hand driving independently of a convention for right hand driving).

Two examples of constitutive convention: a) the already familiar case of words; and b) the case of chess. We have already seen how words have no existence independently of convergences in behavior. The fact that “red” refers to redness is *constituted* by the coordinated use of the sound/inscription “red” in order to refer redness. Sounds and inscriptions have no power, in and of themselves, to refer anything, independently of the coordinated behavior of language users. Now consider the case of chess: there is no context where people move chess pieces about on a chess board, before there is convergence on the rules of chess, and then decide among them that it would be good to have some conventions to regulate the prior activity of playing chess. “Chess” does not name an activity that people can pursue independently of there being rules of chess. To play chess is to follow the rules of chess. The rules of chess are *constitutive* of chess playing.

Constitutive conventions are in stark contrast with musical conventions. For any musical convention imaginable, one can conceive of *musical* behavior *before* that convention is in place. Furthermore, it is the pre-existent pursuit of *musical* goals that makes the conventions intelligible: equal temperament became conventional *because* people had the already musical goal of being able to modulate between keys without sounding out of tune. Convergence in behavior did not establish the *musical nature* of the goals that motivate the convergence. Music incorporates layers of convention, and many musical entities depend, for their existence, on conventions – and, to the extent that those conventions could not exist without language, many musical entities not only depend on language for their existence, but their appropriate understanding *as the kinds of musical entities they are* is not possible without musical appreciation itself being informed by language. At the same time, the contrast with language could not be more striking. If I am correct, both music and words are defined by artifactual functions, but these functions themselves differ in kind. Something intentionally produced and sonically indistinguishable from a movement in sonata form would already be music if no conventions for sonata form were in place; by contrast, no “sound structure” could be a word without social convergence in the use of sounds to represent the world. Convergence makes language but not music, even if music is ultimately dependent on the same mental capacities that allow us to have language.

And this is why I said that the difference between the kind of relational properties in virtue of which sounds (and inscriptions) become words, and the kind of relational properties in virtue of which sounds become music is the basis of the difference between aesthetic theories of art, on the one hand, and institutional theories of art, on the other. What defines an institutional theory of art as *institutional* is that it conceives of art as something that depends on a *constitutive* convention (a convention which is constitutive of the activity it regulates, in the same way that the rules of chess are constitutive of chess), rather than an activity that incorporates regulative conventions but can exist (or whose existence is conceivable) in some form without those conventions; in other words, as something that, like language, depends on convergence of behavior, not only causally, in the way that the existence of an orchard or a building may be causally dependent on coordinated human action, but logically or conceptually.

1.2 Culturally modulated natural-kind activities

This distinction between mere causal dependence on convergent behavior and logical or conceptual dependence is a crucial one.³⁰ It traces back to another distinction I made earlier: all conventions are convergences in behavior, but not all convergences in behavior are conventions. We can understand this distinction more clearly by considering a proposal for classifying theories of art, presented by George Dickie in his article, “Art: Function or Procedure, Nature or Culture”. Here, Dickie divides theories of art into *natural-kind theories of art* (NKTA) and *cultural-kind theories of art* (CKTA). These notions will prove to be immensely helpful. Here is how he defines both types:

NKTA: “A natural-kind theory of art would be one in which it is claimed that art first emerged as a result of natural-kind activity and that art has continued to be created as a result of natural-kind behavior. Of course, this kind of theory can accommodate cultural aspects in the art-creating process – for example, the cultural phenomenon of painting in the impressionist style – but it would maintain that the creative process itself is at bottom a natural-kind activity.”³¹

CKTA: “The institutional theory of art, in either its earlier or its later version, is clearly a cultural-kind theory because it takes a cultural, institutional structure to be the necessary and sufficient matrix for works of art. [...] For the institutional theory, various natural-kind activities may show up in various artworks, but there is no reason to think that any one natural-kind activity is or needs to be present in every artwork.”³²

By “natural-kind activities” (NKA) and “natural-kind behavior” Dickie means those things that are spontaneously done by living organisms; activities like “gathering food, stalking prey, eating, mating, building nests, constructing the elaborate courtship bowers that birds do, living solitarily and living in social groups”.³³ Cultural-kind activities (CKA) and cultural-kind behavior are characterized by being “in some sense invented by the members of a particular group and are passed

³⁰ The same distinction is emphasized by Amie Thomasson in her article “Social Entities”, in LePoidevin et al. (eds.), 2009: p. 546.

³¹ Dickie (1997: p. 26).

³² *Ibid.*, pp. 27-28.

³³ *Ibid.*, p. 25.

on by learning.”³⁴ They are “particular ways of living together, particular ways of hunting, particular ways of raising food, rituals of eating and marriage”³⁵, etc. “Some cultural-kind activities are particular ways that, in one way or another, human beings have come to organize their natural-kind activities.”³⁶

The distinction between NKA and CKA fits neatly into our previous distinction between those convergences in behavior that are conventions and those that are not. All human beings, irrespective of cultural background, eat, sleep, express emotions, etc. These convergences in behavior are not conventions; yet *particular ways* in which these activities are carried out can be conventional (I say *can be* because, evidently, many animals have particular ways of performing activities, without it being the case that such activities are conventional: being a *particular way* of doing something is not a sufficient condition for being conventional, even if it is a necessary condition). What else is necessary? A way of doing something is conventional if it is the result of convergence in behavior (as opposed to being biologically predetermined) and if convergence in behavior could have resulted in a different way of doing the same thing (which is to say that we can imagine the same thing being done differently). This picture also allows us to see the *continuity* between NKA and CKA: at least some of the latter are *modulations* of the former.

Now the question arises as to whether all CKA have a conventional element. Dickie answers in the negative: “An example of a cultural-kind activity that does not involve convention is painting. If one creates a painting, one does so by putting some kind of paint on a surface. This is not to say that painting cannot involve conventions; symbols in painting involve conventions, and depictions may involve conventions.”³⁷

As it is phrased, this seems self-contradictory. Painting is given as an example of CKA that does not involve convention, while the last sentence denies that painting cannot involve convention. However, what Dickie means is something very close to what I said before about music in contrast with words: painting does not essentially involve convention, in the sense that for an action to count as the production of a painting no convergence in behavior is required, in the way that it is required for the existence of words. To put it in other terms: the painterly nature of painting is not a consequence of convergence in behavior. Nevertheless, painting (like music) incorporates layers of

³⁴ *Ibid.*

³⁵ *Ibid.*

³⁶ *Ibid.*

³⁷ *Ibid.*, p.26.

convention, in the use of symbolism and stylistic elements.³⁸ Painting is always done in a particular way, and even the exercise of artistic freedom requires a familiar background in order to be visible as such; yet for any painting convention imaginable one can conceive of painting in the absence of that particular convention. In other words, the conventions of painting are regulative, not constitutive. In that aspect, painting resembles music, and differs from both language and chess.

Of course, someone for whom all representation (including pictorial depiction) is conventional, like Nelson Goodman, would disagree. (This is probably what Dickie has in mind when he concedes that “depictions may involve conventions”, since they obviously can: styles of depiction often are conventional tools, but in that case the conventional element pertains to style, not to depiction as such.) For Goodman, pictures, as well as musical works, belong to symbol systems and as such they are as thoroughly conventional as words. In such a view, depiction differs from linguistic representation, not in what concerns conventionality, but in being syntactically and semantically dense (meaning, respectively, that there is an indefinite number of symbols in a picture and that there is an indefinite number of objects those symbols can denote), and relatively replete (meaning that every feature of pictorial design can have a distinct significance). Other authors propose what are sometimes called “perception” theories of depiction, conceiving it as different in kind from linguistic representation. This is the case with Richard Wollheim, whose theory of “seeing in” parallels in many aspects Scruton’s theory of musical experience. According to Wollheim, what makes a picture represent (or depict) its subject (in other words: what makes a picture the picture it is) is: a) the picture’s author’s intention that it represents X (rather than Y), b) our ability to *see X in the picture* (hence the qualification “perception” theory)³⁹. (Of course, not all pictorial representation works like this, since not all pictorial representation is *depiction*.) “Seeing in” is not a matter of convention, even if it can be inflected by convention: surely, what one is able to see in a picture depends in part on one’s cultural background. Seeing things in pictures is similar to seeing things in clouds, except that in the latter case there are no standards of correctness, no normativity. If two people look at the same cloud and one of them sees a bunny while the other sees a dog, there is no standard according to which one of them is wrong.⁴⁰ Pictures, on the other hand,

³⁸ Cf. Abell (2005: p. 188): “Since more than one arrangement of marks can resemble an object, or be experienced as doing so, etc., particular ways of depicting an object may be conventional even if depiction involves non-conventional relations.”

³⁹ According to Wollheim, the experience of *seeing-in* is twofold: in it we are simultaneously aware of seeing a set of marks on a surface *and* of seeing the depicted object *in* those marks. Scruton will argue that the experience of hearing an organization of tones *in* the sounds is also twofold in this way.

⁴⁰ Here we see how convention may constrain perception: a “thick” cultural background widens the scope of one’s *seeing-in* experiences. My ability to see Pegasus in a cloud or a picture depends on my acquaintance with Greek mythology.

have standards of correctness that derive from a picture's history of production, namely, the intentions of whoever made the picture (in other words, standards of correctness are tied to the artifactual status of pictures). The core idea is not that depiction is always free of conventional elements (which is an obvious falsehood), but that there is no conceptual dependence, as there is in the case of words, between depiction and convergence in behavior (such dependence would have to occur if, as the conventionalist believes, *denotation* is the way pictures represent)⁴¹. In terms of Dickie's classification (and assuming conventionalism about depiction to be false), depiction is an example of CKA that does not (essentially) involve convention, while *seeing in* is a NKA that depiction essentially involves;⁴² particular styles of depiction (say, the tradition of byzantine mosaics, nineteenth century European impressionist painting or Japanese *rinpa*) are CKA: *modulations* of what is basically a NKA, via convergence in behavior.

This point about some CKA being modulations of NKA is a conceptual, not an empirical one. It is an empirical question, e.g., whether or not children's drawings at an early age are culturally inflected. If they are, this counts in favor of the idea that there are no pure (non-modulated) cross-cultural manifestations of depiction. However, this is compatible with the "perception theory" of depiction, according to which "at the core" of depiction is basically a NKA such as the experience of *seeing in*. Whether depiction and *seeing in* are basically NKA is compatible with *all* their particular manifestations being CKA-modulated. The point of classifying depiction and *seeing in* as NKA is precisely that cultural modulation is not *constitutive* of those phenomena, in the same way that being colored is not constitutive of, say, being a flower (even if it's true that all flowers are colored). We can see this by way of analogy: birdsong is not a CKA (in fact, birdsong is not *song* at all: we call it "song" because it is music-like, unlike the patterns of geese squawking or cat meowing), even where birds introduce variations in their "songs". Birds don't build traditions of song making based on learning from cross-generation musical experimentation. If any activity is NKA, birdsong is. Yet, this is obviously not impugned by the fact that there isn't a cross-species (a non-species inflected) repertoire of bird songs.⁴³ Given the kind of creatures that we are (language

⁴¹ Goodman (1976: p. 5): "A picture that represents – like a passage that describes – an object refers to and, more particularly, *denotes* it. Denotation is the core of representation and is independent of resemblance."

⁴² It must be noted that there are other kinds of perception theory of depiction; Wollheim's *seeing in* theory is just one example. But this is not of crucial importance here. The central idea is that some perceptual ability (a NKA) is at the core of depiction. Other alternatives in play are imaginative seeing; experienced resemblance and recognition theories.

⁴³ Here it could be objected that the analogy is a bad one, since cross-cultural variation and cross-species variation are not on the same foot, namely, cross-cultural variation is intraspecific. But the relevant fact is that "singing" is a behavior shared by approximately 4000 species of songbirds. That all its manifestations are species-inflected (that they are *somehow*-inflected) doesn't conflict with its being natural-kind behavior. The relevant aspect in the analogy is that

speaking primates), it is likely that many of our activities can't be other than culturally modulated; that is, locally shaped by convergences in behavior.

Anyway, even if it turned out that all representation (including depiction) is conventional, that wouldn't affect Dickie's point that painting does not (essentially) involve convention. As he says, "If one creates a painting, one does so by putting some kind of paint on a surface." True, though it's also true that not all instances of putting paint on a surface count as creating a painting in the relevant sense. One can, say, apply a coat of paint on a surface for no other purpose than to protect it against corrosion. A "painting" in the relevant sense has some other kind of purpose – representational, expressive, or aesthetic or a combination of these. Now, if it turned out that all representation is conventional, that would exclude depiction from the set of (culturally modulated) NKA that "may show up in various artworks" (though none of them is necessary nor sufficient for art status), according to Dickie.

This digression through painting and depiction was motivated by the need to clarify Dickie's classification of theories of art into his scheme of natural-kind theories of art (NKTA) and cultural-kind theories of art (CKTA). It will also guide us in the application of the same concepts to music, since the view of music on which I base the definition proposed at the outset has structural similarities with the "perception theory" of depiction. Apart from the issue of whether or not depiction is conventional, the other items on Dickie's list of culturally modulated NKA that may show up in diverse artworks are clearly not conventional in the strong sense (though, as I explained, they may always incorporate layers of convention). These are the expression of emotion and the pursuit of aesthetic experience. Thus, according to this scheme, theories of art as imitation (provided conventionalism about depiction is false), as expression and as the production of aesthetic experience will all be classified as NKTA:

"A theorist, for example, who claims that art is the expression of emotion presents a natural-kind theory of art, for the expression of emotion is a clear example of natural-kind behavior."⁴⁴

"If it is assumed that human beings are naturally attracted to basic aesthetic qualities of the kind Beardsley has in mind – unity, brightness, shininess, intense color, and the like – and seek experiences of them, Beardsley also presents a natural-kind theory of art."⁴⁵

its being natural and shared is conceptually prior to its being somehow-inflected. In other words, if there was only one species of songbird that fact would not be *greater* evidence that birdsong is natural-kind behavior.

⁴⁴ *Op. cit.*, p. 26.

⁴⁵ *Ibid.*

Dickie proposes his classification of theories of art against the backdrop of Stephen Davies' previous classification, in his 1991 book: *Definitions of Art*. Dickie offers his own scheme not as an improvement on Davies' but as a way of throwing a different light on the relations between theories of art. This is also the purpose with which I bring Dickie's scheme to bear on the issue of defining music, and how it relates to the definition of art.

Earlier I suggested that different kinds of relational properties were involved in the ontology of different kinds of artifacts (words and musical entities). Artifacts are defined partly by their functions and functional properties are relational properties. They are relations between the object and some aspect of its causal history. Different kinds of functional properties will vary with what stands as the other relatum within the relation. Assuming that a list of three possible relata exhausts the possibilities, we get a threefold classification of how things acquire functional properties.

- 1) As a consequence of convergence in behavior.
- 2) As a direct consequence of intentions involved in their creation.
- 3) As a result of a natural evolutionary mechanism.⁴⁶

Both (1) and (2) concern the functions of artifacts⁴⁷ – human made functional things – while (3) concerns “natural functions”, the functions of biological organs, such as the function of leaves to perform photosynthesis or the function of the heart to pump blood. In all three cases, functional properties are tied to the object's history or its origins. Artifactual functions trace back to *agents*; for this reason they can be called “agentive” functions.⁴⁸ These are the kinds of functions involved in (1) and (2) above. Natural functions are discovered, not imposed by agents; for this reason they can be called “non-agentive functions”. Agentive functions are imposed, either intentionally or as the unintended result of certain intentional actions (convergences in behavior that constitute solutions to coordination problems). Functions beget *norms*: functional entities belong to *normative kinds*. Normative kinds are those that can have well-formed and malformed or dysfunctional instances. A rock cannot “malfunction” as a rock, though it can be more or less appropriate as a paperweight. The norms for paperweights and other functional things are relative to the

⁴⁶ I base this classification on one provided by Zangwill (2014: p. 143). He distinguishes between 1) things with intentionally generated functions, 2) things with unintended emergent functions, and 3) things with only natural functions.

⁴⁷ This includes not only the most familiar kinds of artifacts – physical tools and implements – but also things such as words, concepts and *institutions*.

⁴⁸ To my knowledge, the term “agentive functions” was introduced by John Searle (1995: p. 20; 2010: pp. 58-60). Agentive functions divide in *causal-agentive* and *status-functional*. Such notions will be clarified later on.

aims of agents.

One crucial distinction to be made is between those functions that cannot be performed (and no object conceived as having one such function) independently of a convergence in behavior (particularly, of a *shared representation* that a given object has the relevant function or is enabled to perform it) on the one hand, and, on the other, those functions that can be performed regardless of convergence in behavior or shared representations.

The most common examples of artifacts that come to mind – screwdrivers, paintbrushes, clocks, coffee-makers, etc. – have functions of the latter kind. People don't have to converge in a representation of the coffee-maker as a coffee-maker for it to be able to make coffee; of a clock as a clock for it to be able to give information about the time; of a screwdriver as a screwdriver for it to be able to twist screws; or of a paintbrush as a paintbrush for it to be able to spread paint on a surface. The production of those objects may depend on joint action, but their functions are not conceptually dependent on convergent behavior or shared representations. Precisely the opposite is the case with words: a word exists in a language when people jointly recognize some sounds or squiggles as having (or being enabled to perform) symbolic functions, i.e., when a set of sonic/graphic patterns⁴⁹ is used as a *norm kind* with symbolic functions. The idea that private languages are impossible expresses the same difference I'm pointing out between the kind of relational properties in virtue of which artifacts of the former kind are the artifacts they are, and the kind of relational properties in virtue of which artifacts such as words are what they are. All functional things derive their functions from some of their *historical* properties: the heart derives its function of pumping blood from a causal mechanism; most words derive their functions from convergences in behavior that constitute solutions to coordination problems;⁵⁰ coffee-makers, clocks and paintbrushes derive their functions from *intentions* that a certain physical structure performs the relevant function, intentions which (in order to be the relevant ones) must be connected to an *insight* according to which a certain kind of object will have the causal powers or dispositions appropriate to "implement" the function. Convergence in the *use* of an artifact issuing from such intentions and insights is an extra feature, irrelevant to its status as the functional thing it is.

Imagine that instead of convergence in the use of equal temperament, European composers had settled for modal music and static harmonies until the twenty-first century; suppose that musicians

⁴⁹ I say a *set* of sound/graphic patterns since *the same word* can have a variety of spellings and pronunciations. The norm kind comprises all spellings and pronunciations (abstract patterns) that are *recognized* as correct.

⁵⁰ Though Ruth Millikan (1998: p. 162) has called attention, *contra* Lewis, to the fact that not all conventions arise from coordination problems. In the case of language, she gives the example of expletives.

simply did not care about modulation between keys, eschewing functional harmony the same way it is said that ancient Egyptian iconography eschewed originality. Still, in such a scenario it is perfectly conceivable that some “underground” or *avant-garde* musicians would write *unconventional* pieces for a system of equal temperament (the same way a restricted group of musicians did create twelve-tone serial pieces in a world dominated by major/minor diatonic tonality). Actually, the introduction of equal temperament and functional harmony required a vast amount of cooperative, unplanned experimentation (just as the appearance of twelve-tone serialism required vast amounts of experimentation with the boundaries of western tonality); but this is no more relevant than the fact that the making of a cathedral requires vast amounts of cooperative action is relevant for the *status* of the cathedral as a building (as the kind of artifact that it is). If someone endowed with superhuman abilities single-handedly built a cathedral, it would still be a cathedral; but superhuman abilities would be of no avail in the attempt to create a private language, or in making it the case that some sound/squiggle type counts as a word, in the absence of convergent behavior. Likewise, there is nothing *incoherent* in the idea of a thirteenth century composer endowed with superhuman intelligence, discovering equal temperament on his own and creating twelve-tone serial pieces. Such a composer may be a physical or natural impossibility but certainly not a logical or metaphysical one.

Our twelve-tone equal temperament is just one in a myriad ways of dividing the octave, an endeavor with which human beings have been tinkering and experimenting at least since ancient Greek antiquity. It plays with something which is given in nature (the pitch continuum and the universal human characteristic of perceiving pitches an octave apart as the same though qualitatively distinct pitch), but it is not itself naturally given (since, more than a slicing up of the continuum, it is a *tool*; away of *organizing* pitched sounds so that they can “do things” they can’t do, or not as well, when organized differently – which of course is only “good” to the extent that there is an interest in pitched sounds doing those things). It had to be arrived at. And it only has a functional purpose given the kinds of creatures that we are: mindless nature doesn’t care about modulating freely between keys without sounding “out of tune” (the “heavenly spheres”, unfortunately, are tone deaf!). It is not merely a division of the octave for division’s sake; it has drastic aesthetic repercussions on the way we create “movement”, “atmosphere” and “color” in the music and that is the reason why there is to be an interest in how to divide the octave.

In sum, twelve-tone equal temperament is a human artifact. It is also one solution for some musical problems (e.g. how to make “modulating music”), though not for all; and, like when translating poetry from one language into another, there are things expressible in other tuning systems that can’t be expressed in equal temperament and vice-versa. This is what drove some

twentieth century western composers such as Lou Harrison and Harry Partch to experiment with alternative systems, like just intonation. It differs from coffee-makers and paintbrushes in that it has no spatial location (e.g., it makes no sense to ask where equal temperament is nor to say that it is five miles away to the north) or, if one adopts that brand of realism, it is located wherever an instrument tuned to that temperament is located (the problem with this view is that it was already possible to use equal temperament before any instrument was tuned in accordance). However we choose to describe that feature, it marks a distinction between artifacts: on the one hand, there are artifacts like screwdrivers, coffee-makers, oboes and cellos; on the other hand, there are artifacts like words, the dominant seventh chord, the Dorian mode, and twelve-tone equal temperament. For a matter of convenience I'll adopt the terminology of "concrete artifacts" and "abstract artifacts". By using this terminology I am not necessarily committing myself with a substantial ontological thesis about those things qualified as "abstract". In this way, talk of abstract artifacts is as harmless as talk of properties for a nominalist who has a few paraphrases up his sleeve. One alternative would be to talk of "non-concrete" rather than "abstract" artifacts, though these seem synonymous, if we take concreteness and abstractness to be mutually exclusive. However, "non-concrete" could be taken here to mean that the item in question is not to be identified with any given particular concrete physical thing, short of any commitment to realism about abstract objects.

1.3 Sortals, mind-dependence and real existence

What falls under a humanly invented conceptual category possesses a nominal, not a real essence. All concepts are humanly constructed, but some such as 'jade' are supposed to latch on to natural structures that exist independently of human thought, while others such as 'parking ticket' latch on to artificial structures, functions, intentions, institutions, or practices.

Stephen Davies⁵¹

Some things exist and are how they are independently of our thoughts and actions – things such as rivers, mountains, and trees. Of these, some are what we call “natural kinds” – for instance, water, silver, copper, and more controversially, biological species.⁵² To exist independently of human thought is not the criterion by which we single out natural kinds as such, though all natural kinds exist independently of human thought. In other words, being independent of human thought is a necessary but not sufficient condition of being a natural kind. For instance, mountains exist independently of human thought, that is, the masses of rock to which we refer via the word “mountain” do not, in any manner, depend for their existence on how we think about them or refer to them. But things are not always so straightforward. “Mountain” is not a natural kind term, since what things count as mountains is a fact dependent on our categorization of reality – e.g. whether an island counts as a (partially submerged) mountain depends on how we decide to count mountains, that is, it depends on what restrictions govern the use of the sortal “mountain”. But how do we reconcile this with the idea that the existence of mountains is independent of our thought? Once the conditions governing the application of the sortal are fixed (thus settling issues such as whether certain islands count as mountains), it is, so to speak, up to the world whether *it* contains objects that fit the sortal by satisfying the associated conditions. Those things that fit the sortal are what they are independently of how we think about them. For instance, masses of rock that fit the sortal “mountain” have the physical features they have independently of there being a sortal that they fit in virtue of satisfying conditions for its application, conditions which specify the relevant features possessed by the bodies of rocky mass to which the sortal is to apply. Our way of conceptualizing

⁵¹ Davies (2010: p. 25).

⁵² It is crucial not to confuse natural kinds with natural kind *terms* and natural kind *concepts* expressed by those terms. So, the natural kind term “water” expresses the natural kind concept WATER and it refers the natural kind water, i.e., the kind of stuff that fills our lakes and oceans, falls from the sky as rain and runs from taps.

reality, counting things using sortals (the only way there is of counting things)⁵³, is powerless to determine whether anything *actually* fits the conditions fixed for the application of some of our concepts.

What depends on our conceptualization, then, is not whether mountains *exist* (whether the world contains things that fit that sortal) and whether they have the intrinsic and extrinsic features they have, but whether some of our concepts apply to independently existing things (such as masses of rock). So to assert that mountains have an objective, mind-independent existence is perfectly compatible with asserting that whether *X* counts as a mountain depends on our way of classifying things. At bottom, what I'm saying is that with terms such as "mountain", i.e., terms that do not name natural kinds, the status of borderline cases ultimately depends on what restrictions we *decide* to associate with the term. This was what Morris Weitz dubbed the "open-endedness" of empirical concepts: the feature in virtue of which applications of the concept to borderline cases depend on a *decision* about whether or not to extend the concept to accommodate the said case.

It is important to emphasize that the status of a concept as "open-ended" doesn't depend on there being *actual* borderline cases that call for a decision, but merely on the possibility of such a circumstance. Here is how Weitz characterizes open concepts:

"A concept is open if its conditions of application are emendable and corrigible; i.e., if a situation or case can be imagined or secured which would call for some sort of decision on our part to extend the use of the concept to cover this, or to close the concept and invent a new one to deal with the new case and its new property."⁵⁴

In short, a concept is open-ended if it *possibly* has genuine borderline cases. So that even if for a given concept there are no *actual* borderline cases (as once apparently there were no counterexamples to the idea that art is imitative or representational), the concept is still open-ended if it has a counterpart (in some possible circumstance) with which is only partially co-extensive (if there is no overlapping at all, then we have merely two homonymous but utterly distinct concepts). This makes it unclear whether or not *all* empirical concepts are open-ended, but that need not worry us here.

⁵³ Zemach (1986: p. 241): "The question, how many objects are there (say, in the forest) makes no sense at all. To count, one needs a count noun (a sortal); one may then count trees, or leaves, or molecules, or what have you, but not things *simpliciter*. There are identity conditions for trees, for logs, and for clumps of wood, but there are no identity conditions for things." – The apparent counting of things without sortals, as in " $2 + 2 = 4$ ", is actually shorthand for "two of *any sort* plus two of *any sort* equals four".

⁵⁴ Weitz (1956: p. 31).

There seems to be an asymmetry between terms like “water” on the one hand, and terms like “mountain” on the other, in that in the former case there is a property which all and only things falling under the concept instantiate, i.e., the property of being H₂O, while in the latter none such property or set of properties is forthcoming. In other words, things in the extension of “water” (and other natural kind terms) seem to have clear conditions of individuation, which contrasts with the case of “mountain”, since what counts as a mountain depends in part, as we’ve seen, on what people *decide* (on what is explicitly or implicitly agreed by a community of language users) to count as a mountain.

At first this doesn’t seem to differ much from the fact that “water” only applies to the material stuff water because there is a convention or a use connecting the word to what it refers. We picked up the concept by immersion in our linguistic community and, using it, we implicitly make use of the restrictions associated with the term (in other words, we learn by deferring to other language users). The restrictions could have been other and, accordingly, the term “water” could have been used to refer a different substance. (As it happens, the restrictions could be qualitatively indistinguishable and yet the term would have a different extension, as is suggested by Twin-Earth scenarios.) But this is relevantly different from the case of “mountain”, where a degree of vagueness is built in the concept. In a Twin-Earth scenario, WATER could become a disjunctive concept, like the concept JADE. The relevant difference is that in either the Twin-Earth scenario or the actual case of “jade”, the disjunctiveness of the concept has no bearing on the fact that the things referred to have clear conditions of individuation. It just happens that the concept is ambiguous between two natural kinds. Conceptual disjunctiveness affects our thought and talk about reality, not reality itself, though it raises the question of the justification or usefulness of grouping the disjuncts under a single concept. In the case of “jade”, the justification hinges on a relation that holds between us and the things referred to via the disjunctive concept, namely, our valuing of certain perceivable properties possessed by the things referred, though other considerations are also involved (we don’t just classify as “jade” *anything* that shares the perceivable properties of jadeite or nephrite). Both in the case of “water” and “jade”, the fact that both the single substance referred by the former and the two substances referred by the latter have clear conditions of individuation does not depend on any decisions taken by language users regarding the use of the concept. The fact that there is one essential individuating property for water, and that “jade” refers to either of two substances each with essential individuating properties, was not settled by our linguistic practices but was the object of a *discovery* that occurred well after the linguistic practices were firmly established. And here lies the difference between the concepts WATER and JADE, on the one hand, and the concept MOUNTAIN on the other: in the case of the latter,

no such discovery seems available or even possible. Mountains have no “hidden essence” that we could discover, independently of our firmly established practice of classifying certain masses of rock as “mountains”.

The relevant similarity between all three cases, though, is that independently of issues raised by conceptual vagueness, disjunctiveness and open-endedness in the face of possible circumstances requiring a revision of the conditions of application, the *source* of those conditions associated with the concepts seems undoubtedly located in the external world, independently of what we think. Even though mountains have no language-independent essences and the classificatory practice in this case seems to hinge on a relation between the external objects and the interest we take in them, the constraints that make the concept intelligible are firmly anchored in features of the external world which are independent of the way we represent the world. Our interest in such features is not ontologically constitutive of the features themselves. These were already available “out there”, before our interests started to play their role in how we divide the world into categories of things. In Stephen Davies’ words, some of our concepts seem to “latch on to natural structures”⁵⁵, which exist and have the properties they have independently of those concepts and of issues such as vagueness, open-endedness and disjunctiveness, which at least in the cases we have considered, seem to pertain only to how we conceptualize an independently existing reality. The point is that the (ontological) *objectivity* of the natural structures to which our classificatory practices “latch on” seems unaffected by issues such as whether an essentialist characterization for the extension of a given concept is available. In some cases, empirical discovery and conceptual fine-tuning will render such characterizations available, in other cases it won’t. That doesn’t obliterate the fact that certain mind-independent necessary conditions for the application of sortals must obtain; e.g., in a non-figurative, non-extended use of the sortal “mountain”, nothing can fall under the concept unless it is a body of rock mass produced by certain geological phenomena (though this is not a *sufficient* condition).

In stark contrast with rivers, mountains, water and silver, there is a class of entities which seem to pose more serious problems. These are termed *social kinds* and *artifactual kinds*; categories of things whose existence is *not* independent from what (human) agents think and do. In the case of such categories, it’s not that we have concepts that latch on to independently existing structures but rather the “structures” referred to with social and artifactual kind concepts have no reality independently of human agents, their interests or convergences in behavior. Artifactual kinds include causal-agentive artifacts (artifacts whose functions derive from intentions and can be

⁵⁵ S. Davies (2010: p. 25).

performed independently of convergence in behavior) and status-functional artifacts (artifacts whose functions derive from a convergence in behavior and can either be intentional or unintended). Social kinds include social practices, conventions and institutions. Both categories partly overlap, since all status-functional artifacts are institutional artifacts.

How are we to conceive of the objectivity of social and artifactual kinds given that these, unlike things referred to with those concepts that “latch on to natural structures” (which include natural kind concepts and “culturally contaminated” concepts like mountain), don’t exist independently of human thought and action? How can there be an objective reality, a domain of objects, facts, events, that partly depend on our propositional attitudes? If something is a fact partly because there is a shared belief that it is a fact or because it stands in a causal chain that traces back to the intentions of agents, won’t that be a merely subjective or, in other terms, a “socially constructed” fact?

In trying to provide an answer to this question, John Searle (1995: pp. 7-8) distinguished what he calls the “ontological sense” of the contrast between *objective* and *subjective* from the “epistemic sense” of this very same contrast. In the ontological sense, “objective” and “subjective” qualify *facts* and not, as in their primary use, *judgments* (attempts at stating the facts). Accordingly, ontologically objective facts are those facts that are independent of our experiences, thoughts or propositional attitudes (e.g. the fact that the world contains mountains).⁵⁶ Ontologically subjective facts, on the other hand, are those that somehow depend on cognitive agents (e.g. the fact that there are beliefs about the existence of mountains). In the epistemic sense, “objective” and “subjective” do not qualify facts but *judgments*, i.e., purported *statements of facts*. Thus, there are epistemically objective judgments, i.e., attempts at stating the facts (where the purportedly stated facts are either objective or subjective in the “ontological sense”) which can be (objectively) true or false (e.g. the statement that Vaughan Williams wrote his oboe concerto in 1944). Epistemically subjective judgments are those that either have no claim to normativity or whose normative claims are groundless, i.e., statements whose truth or falsity can’t be objectively settled, because they depend on feelings, attitudes or points of view. Searle’s preferred example here are judgments of aesthetic merit (though the topic of whether there is such a thing as aesthetic normativity is the object of controversy) or comparative artistic merit (e.g. the judgment that Rembrandt is a better artist than Rubens). Whether or not aesthetic judgments turn out to have justified normative claims, this nevertheless provides the right kind of example: *if* aesthetic judgments lack normativity, then they

⁵⁶ So, under the present classification, the fact that there are mountains is an ontologically objective fact; the fact that the sortal “mountain” refers to mountains is an ontologically subjective fact (nothing counts as a “sortal” or functions as such independently of any cognitive agents’ thoughts and beliefs).

are epistemically subjective, in Searle's sense of not "corresponding" to any objective aesthetic *facts* in the world. Notice, however, that under this classification, the fact that people have aesthetic *responses* falls under the notion of "ontologically subjective fact", about which there can be epistemically objective statements. For instance, *if* aesthetic judgments lack normativity, *then* it is an objectively true statement about people's aesthetic responses that nothing about them grounds the normative claims of aesthetic judgments. This assertion of aesthetic subjectivism, as well as the assertion of any rival theory about the nature of the aesthetic, counts as an epistemically objective statement about an "ontologically subjective domain": either one assertion is objectively true or all are objectively false.

It strikes me as a consequence of the previous scheme that the so called "ontologically subjective facts" form a subclass within the wider class of "ontologically objective facts". What this means is that the contrast between objective and subjective is not to be understood in terms of mutual exclusion, as one might construe, e.g., the contrast between *concrete* and *abstract* (in the sense that something's being abstract excludes its being concrete and vice-versa). Maybe ontologically subjective facts are better understood as a special sort of relational facts *within* ontologically objective facts. In that case, "ontologically subjective" and "ontologically objective" would not name two mutually exclusive categories of facts but rather two levels of description for the same kind of fact.

Take for instance the fact that someone experiences a certain sequence of sounds as delicate and graceful. This is an ontologically subjective fact, since delicacy and gracefulness, even if thought of as qualities of objects, are only intelligible in relation to the experiences of cognitive agents; but this very same fact – that the qualities of delicacy and gracefulness must involve reference to an experiencing subject – is in its turn an ontologically objective fact: independently of what we think about it, the world is such that the qualities described by the predicates "graceful" and "delicate" essentially involve the experiences of cognitive agents. Whether *ascriptions* of delicacy and gracefulness (as well as other aesthetic qualities) are (epistemically) objective or subjective (whether they express *propositions*, which are either true or false) is a separate matter. Either way, the fact that we use aesthetic predicates to describe certain experienced features of objects, the fact that we have such experiences, is itself as objective as any other relational fact that doesn't involve a necessary reference to human psychology or the experiences of cognitive agents. Thus, in the "ontological sense" *everything* is perfectly objective (in one level of description), including the "subjective" aspects of reality – like a listener's experience of a sequence of sounds as delicate and graceful. In the "epistemic sense", the difference between objective and subjective lies in the normative claim of the *statements* we make about how things stand in the world, independently of

whether or not these involve a necessary reference to the experiences, thoughts, and propositional attitudes of cognitive agents. For instance, the assertion that a certain sequence of sounds is “delicate and graceful” will be subjective in this sense if it’s just “a matter of taste” or of an individual response to the object, if there is no way to evaluate the responses of individuals against an objective standard (i.e., if *delicate* and *graceful* is just what is delicate and graceful *to me* or any other subject, which begs the question of how I or anyone else came to master the use of aesthetic predicates in such a way that it accords with how they are used in the community of language users).

It is important to note that Searle’s distinction between the “ontological sense” and the “epistemic sense” of the objective/subjective contrast apparently makes no room for subjective judgments or statements (e.g. those which are based on hedonic responses) with a claim to normativity. It merely assumes that they have no such claim. Maybe they don’t, but one must bear in mind that there is a tradition in philosophical aesthetics which acknowledges the subjective character of a certain class of judgments (their reliance on hedonic responses), namely, *aesthetic judgments*,⁵⁷ but sees this as perfectly compatible with a “standard of taste” or some source of normativity against which to evaluate such judgments. But I don’t want to get into epistemological hurdles of this kind here. The point of Searle’s distinction is simply that the ontological subjectivity of a given domain is not incompatible with an epistemically objective *discourse* about that domain. As far as aesthetics is concerned within these distinctions, it suffices for my purpose to establish with some degree of plausibility that the *reality* of a domain we can appropriately describe as “aesthetic” (a domain of experiences, interests, responses, valuations) is independent from the epistemological issue of whether *ascriptions* of aesthetic properties, value or merit can be said to have justified normative claims. Yet another issue is whether what we call “the aesthetic” can itself be appropriately described as a social kind, or something that is “socially constructed”, over and above the fact that aesthetic *talk* is evidently socially constructed: aesthetic predicates are linguistic entities and we learn how to use them correctly by immersion in a linguistic community (we need language in order to make judgments, and the fact that language is “socially constructed” is compatible with some of the judgments we make being epistemically objective, for instance, the judgment that language is “socially constructed”).

The main worry about social and artifactual kinds is whether their ontological subjectivity (i.e., their dependence on the actions, shared beliefs and intentions of human agents) constitutes an

⁵⁷ More precisely, judgments of *taste* within the class of aesthetic judgments, which also included “judgments of the agreeable”, the difference being that the latter, unlike the former, have no claim to normativity.

obstacle to there being an epistemically objective account of social and artifactual kinds, the same way we can have an epistemically objective account of those “natural structures” that some of our concepts “latch on to” (whether these are natural kinds or “natural structures” specified by “culturally inflected concepts”, like mountain).

Why is this relevant? As we saw, a characteristic shared by those things referred to by terms like “water” or “mountain” is that their existence is independent of how we pick them out, how we refer to them. There is a clear distinction here between conceptualization and reality, between the way we capture such “natural structures” in concepts or linguistic descriptions and the brute fact of their existence and their having the intrinsic properties they have, independently of our picking out those properties via words and concepts. But where no such clear dividing line between conceptualization and reality is forthcoming, as would seem to be the case with social and artifactual kinds, given their conceptual dependence on human thoughts and behavior, there is room to doubt whether an epistemically objective account of their nature (i.e. an account whose truth is independent of what people believe about the subject) is possible.

In other words, if something is “socially constructed” (as social kinds by definition are) or somehow depends on the intentions of agents, can it nevertheless have an objective nature, in the sense that it is subject to certain logical and metaphysical constraints (conditions something must satisfy for it to be *of* the relevant kind), and is our discourse about such things capable of being true or false? Furthermore, are objective truths about social and artifactual kinds *discoverable* or do cognitive agents have some sort of privileged epistemic access to them, in the sense that they can’t be fundamentally mistaken about the nature of social kinds?⁵⁸

⁵⁸ Amie Thomasson has an argument for metaphysical descriptivism in the ontology of art which draws on the idea that grounders and regrounders of art-kind terms can’t be fundamentally wrong about the nature of the things referred to with such terms since, in her view, disambiguation (the solution of the so-called *qua problem*) depends on the “nascent ontological conception” that (re)grounders of the art-kind term associate with it, thus setting the conditions anything must satisfy for it to be a token of that kind. Julian Dodd criticizes this view by arguing that it depends on a theory of reference by fit, which can be replaced by a more plausible theory (e.g., Evan’s hybrid theory of reference) that doesn’t entail any epistemic privilege for (re)grounders of art-kind terms.

1.4 Disjunction and real definition

It is crucial to note that even concepts which cannot be defined in terms of a set of necessary and sufficient conditions, may well have certain defining features.

Andrei Marmor⁵⁹

The above quote could be rewritten thus: even for those cases where we are unable to pinpoint necessary *and* sufficient conditions for something to be of kind *K*, that does not mean we are not able to pinpoint *either* necessary (but not sufficient) conditions *or* sufficient (but not necessary) conditions for something to be of kind *K*. If a feature is either necessary or sufficient for something to be of kind *K*, then that feature is a *defining* feature of *K*.

A specification of necessary (but not sufficient) or sufficient (but not necessary) conditions for something to be of kind *K* counts as a definition of *K*, but not as a *real* definition (i.e. a definition that pinpoints conditions which are both necessary and sufficient for something to be of kind *K*). Definitions in terms of necessary (but not sufficient) features will be too *wide* (in the sense that they capture all things which are *K* but also some things which are not *K*) and definitions in terms of sufficient (but not necessary) conditions will be too *narrow* (in the sense that while they do not capture things that are not *K*, neither do they capture all things that are *K*). Such definitions, nevertheless, can be useful *characterizations* of the kind in question. For instance, the definition of music as “organized sound”, if true, is informative about the nature of music, though it also captures things which are not music. The property of being a symphonic movement written in sonata-form is sufficient for something to be music, but this is too narrow, since most things which are or ever were music are not symphonic movements written in sonata-form. Since necessary conditions for *K* are conditions that *all* things of kind *K* satisfy, a definition in terms of necessary (but not sufficient) conditions for *K* is more informative than a definition in terms of sufficient (but not necessary) conditions. However, a list of sufficient (but not necessary) conditions for something to be of kind *K* *can* be more or less illuminating about the nature of *K*, depending on how “central” these conditions are. A conjunction of properties (or conditions) which, taken individually, are neither necessary nor sufficient, may nevertheless be sufficient (even if not necessary) for *K*. And different combinations of such properties can also be sufficient though not necessary for *K*.

⁵⁹ Marmor (1996: pp. 350-351).

Though the term “definition” (of art, music, etc.) is often used interchangeably with “theory of” (art, music, etc.), there’s something to be said for the careful distinction of both. One and the same definition of an art kind *K* may be compatible with more than one theory of that art kind *K*. For instance, the definition of music with which I am here presenting – music is the art of tonal sound – will be compatible with more than one account of what it is to hear sounds as tones, to hear movement, action and emotional qualities in mere sequences of sounds. Scruton’s “metaphorical transfer” view of musical experience is but one such account. Though Scruton never offered a definition of music, the above definition could be combined with his account of musical experience.

The goal of a definition of an art kind *K* is to provide *defining features* of *K* (necessary, sufficient or both), while the goal of a theory of an art kind *K* is wider, since there will be more to an art kind than a list of essential properties⁶⁰. For instance, we may wish to know why is there an interest for us in singling out things that are *K*, what value they might hold for us. A theory of *K* may or may not include a definition of *K*, whether a *real* definition (in terms of necessary and sufficient conditions) or a *characterization* (in terms of necessary but not sufficient, or sufficient but not necessary conditions).

The expression “definition of (*K*)” is also often taken as short for “definition of *the concept of* (*K*)”. However, we should carefully distinguish both notions, for the following reason, among others. The use of concepts is multifarious and subject to change over time, so that the same concept may specify or refer different kinds at different times (for instance, the concept *meat* originally referred to any edible item, and later it came to refer *flesh* only)⁶¹ and under different uses (secondary meanings and metaphorical uses of literal meanings being the most straightforward examples). This means that, depending on its *use*, one concept may actually refer to different kinds of things. And among the different things referred to by a single concept there may be an interestingly unified range of objects or events. As such, we should beware of the ambiguity in “definition”, between defining concepts (constituents of thoughts about items in the world) and defining kinds (a unified range of items in the world). The ambiguity is between two distinct types of definition: to define a concept is to provide a nominal definition, i.e., it is the same as giving the meaning of a word; to define a kind *K* is to specify the property or properties that unify a range of items in the world (the conditions anything must satisfy to be of kind *K*), independently of our use of concepts or words. When speaking of definition, my primary concern is with kinds, not concepts, and thus I am not primarily interested in the concepts/words *art* and *music*, but in whether among

⁶⁰ See S. Davies (2010), chapter 1, section III.

⁶¹ I took this example from Sainsbury & Tye (2012: p. 46).

the things referred to with those concepts/words there are interestingly unified kinds and what unifies them.

Not all (putative) real definitions of a kind *K* have the classical form of a list of one or more conditions, individually necessary and jointly sufficient, for something to be an instance of *K*, though it is a matter of dispute whether any definitions lacking the classical form can be appropriately labeled “real definitions”. Other forms of definition are *disjunctive* and *cluster* definitions, which are a slightly more complex sort of disjunctive definition. One relevant distinction is that simple disjunctive definitions are not typically open-ended. It is the open-endedness of cluster definitions that casts an extra doubt on whether they can count as real definitions.

A well-known example of disjunctive definition is that of jade, which actually comprises two distinct substances: jadeite and nephrite; so that whatever is jade is either jadeite or nephrite. Neither condition is necessary but either one is sufficient for being jade. It can be argued, though, that the disjunction of properties yields a necessary condition: it is necessary, in order for something to be a *K*, that it be either *A* or *B*. Whether disjunction makes a definition dysfunctional is the object of dispute.⁶²

There are quite a few examples of disjunctive definitions in the literature on art. Thus, for Beardsley, “an artwork is *either* an arrangement of conditions intended to be capable of affording an experience with marked aesthetic character *or* (incidentally) an arrangement belonging to a class or type of arrangements that is typically intended to have this capacity”,⁶³ and for Tatarkiewicz, “an artwork is either a reproduction of things, or a construction of forms, or an expression of experiences such that it is capable of evoking delight, or emotion, or shock”.⁶⁴

A cluster definition assumes the following form: something is of a kind *K* only if it satisfies a minimum cluster of conditions, from a specified list, none of which are individually necessary for *K*-hood. The result will be a disjunction of the several possible combinations of (non-necessary) conditions, the number of which will depend on how extensive is the list from which the combinations or clusters of conditions are made. So, for instance, if something is art just in case it satisfies at least eight out of ten conditions, there will be fifty six possible combinations or ways for

⁶² See J. Kingsbury & J. McKeown Green (2009).

⁶³ Beardsley (1982: p. 299).

⁶⁴ Tatarkiewicz: (1971: p. 150).

something to be art⁶⁵. Some proponents of the cluster account of art, like Berys Gaut (2005), do not present it as a real definition but as an anti-essentialist characterization, under the assumption that the list of conditions is always open to future revisions. In this view, a cluster account is a way of strengthening Weitz's notion of art as an "open concept" and his "family resemblance" approach to how we identify items as artworks, whence things falling under the art concept are not bound together by any exhibited property or set of such properties but by "strands of similarities" or "family resemblances": while the list specifies the relevant similarities without which the appeal to resemblances would be vacuous (since any two things will resemble each other in some way), the possibility of there being cases which call for a decision whether to expand the list remains open.

Other philosophers, like Denis Dutton, present their own version of the cluster account as a real definition. The argument here is that the potential revisability of the list of conditions bears no weight on whether the cluster account provides an anti-essentialist characterization rather than a real definition. Dutton builds his cluster definition from a list of twelve features that count towards an object being a work of art, should it have a significant number of them. He places the burden on the anti-essentialist to show that something could be art while lacking all twelve features in the list. If no examples can be found, there is no reason to think that the definition, while lengthy and disjunctive, doesn't capture the essential properties of art, even if in principle it is open to revision. Our aesthetic preference for simpler definitions doesn't show the cluster account to be false or inadequate as a definition (even if "classical" definitions are, *ceteris paribus*, preferable to cluster definitions), and to suppose that the revisable character of the list leads us to infinite disjunction is to incur in a slippery-slope fallacy, keenly noted in Dutton's rejoinder: "A thousand or more ways of being art is still a long distance from an infinite number of ways to be art."⁶⁶

Much more can be said about the definitional/non-definitional nature of cluster accounts, but I'll not be too concerned with that here, since I don't think it necessary to adopt a cluster account of music. I will be concerned with whether an aesthetic definition of art must be disjunctive in order to deal with "hard cases", and whether an aesthetic definition of music must mirror that disjunctiveness for the same reason.

As (at least) a characterization of art, the cluster account is highly plausible, since it can combine elements from all definitional approaches without committing to any of them. A cluster theorist may consider that the properties specified by each (or some) competing "theory of art" are not necessary although sufficient, or at least part of a sufficient cluster of conditions for arthood. This

⁶⁵ The example is taken from Stephen Davies's (2004) discussion of Berys Gaut's cluster account of art.

⁶⁶ Dutton (2006: p. 375).

comes with a certain cost, though. As Robert Stecker points out: “To think of art as a cluster concept makes it permanently indeterminate which set of properties, at which weighting, suffice or are needed for an item to be an artwork. This means that hard questions cannot be settled.”⁶⁷ We have no clear way of specifying how minimal a minimally sufficient cluster of properties for art can be, nor do we have a clear way of determining their relative weight or priority – maybe some conditions in the list count more heavily towards an item being art than others. This need not be understood as an embarrassment for the cluster account as such. The vagueness or indeterminateness just pointed out might merely reflect something about the ontology of social kinds, as opposed to the ontology of natural kinds. Furthermore, there might be asymmetries between social kinds, regarding how exacting their metaphysical constraints are. In other words, the epistemological vagueness of the definition might merely reflect the metaphysical vagueness of art itself. If art lies in a continuum with non-art, in the sense that an indefinite range of activities may incorporate art elements without being art *tout court*, this is only to be expected. Art as a distinct activity emerges from the whole of our “aesthetic lives”, in virtue of “natural-kind activities” such as our seeing marked surfaces as depictive representations, our fondness of representations and our allure for “basic aesthetic qualities”.

Dutton isn’t particularly concerned with “hard cases” since he thinks it’s methodologically ill-advised to theorize about the nature of art from borderline or “hard” cases. To understand what art is, we must begin by the undisputed center, not the fringes (2006: p. 375). He draws an analogy with legal concepts: it is bad methodology, while trying to define murder, to begin by discussing euthanasia and abortion. We only approach the hard cases after we have a firm grasp on the central, undisputed ones (*ibid.* p. 368). We certainly don’t assume that a theory is *prima facie* better for including all “hard cases” within the concept’s extension, since we have no reasons to count maximal inclusiveness as an epistemic virtue. When it comes to the philosophy of art and the individual arts, not only the temptation to begin by discussing ready-mades, conceptual art, 4’33” or *musique concrète* seems almost irrepressible, it also seems like the natural place to start, since we feel immediately puzzled by the status of “hard cases” as art in a way we don’t feel puzzled by the status of “central” cases of art (which is why they are “central” in the first place). However, just because a Haydn string quartet is an undisputable case of music (in the sense that if anything is music then a Haydn string quartet is), it doesn’t mean we’re any more enlightened about what makes it music, even if we allow that our common use of the concept is informed or guided by

⁶⁷ Stecker (1997: p. 25).

some form of implicit knowledge of this. It is one thing to be able to recognize an object or event as an instance of art or music (an ability we acquired while becoming competent users of the language), quite another to be able to spell out why we have the need to distinguish a range of items as art or music, in other words, to explain the point of there being concepts for such things in the first place. The fact that we (usually) don't have difficulties in applying the words "art" and "music" in contexts similar to those where we learned how to use the words in the first place makes us no wiser regarding the deeper reasons (if there are any) in virtue of which central cases of art and music are central cases; it makes us no wiser as to whether such things as "paradigms" of art and music have, to use Stephen Davies' expression, more than a "purely nominal essence"⁶⁸, or if our classifying them as such is guided by an intuitive (though defeasible) grasp of an underlying "real nature" shared by paradigmatic cases, in virtue of which the "hard" or borderline cases could then be properly judged as falling (or not) under the concepts ART and MUSIC.

Though I believe cluster definitions to be equally compatible with essentialism and anti-essentialism about art, I see their role as more epistemological than metaphysical, that is, I think they are more useful in specifying identification criteria for art than in describing the essential properties of art. This is not just because they are equally compatible with essentialist (Dutton) and anti-essentialist (Gaut) views on art, but also because they don't commit us to any particular definitional approach. What this means is that the truth of a cluster definition of art is compatible with the truth of an explanatorily more basic definition, whether functionalist or proceduralist.

An example of this can be seen in what is perhaps the most important difference between Gaut's and Dutton's cluster accounts of art, that is, the fact that the latter, unlike the former, doesn't include "possessing positive aesthetic properties" as one among the listed art-making features. The reason for this difference lies in the way Dutton conceives of aesthetic properties:

"To say that art by its nature possesses aesthetic properties has led theorists to beg important questions. (...) I reject this view, and along with it the notion that possessing aesthetic properties is something that might be added to the list. To the contrary, it is the combination of the other items on the list – virtuoso skill, novelty, imaginative representation, emotional feeling, intellectual challenge,

⁶⁸ The idea of a "purely nominal essence" is, roughly, this: things classified as members or elements of a class *K* have no deeper criteria of *K*-membership than the fact that we classify them as such. In other words, they aren't bound by any common, empirically verifiable, "inner essence" in virtue of which they are elements or members of *K*.

and so forth – which, combined in the experience of a work of art, are precisely the aesthetic properties of the work, normally delivering the pleasure described as Item 1.”⁶⁹

This passage can be construed as the specification of a wide *dependence base* for aesthetic properties, that is, the list specifies non-aesthetic properties (both perceptual and contextual, thus the qualification “wide”) on which the aesthetic properties of works depend. In other words, aesthetic properties are not properties an artwork has over and above the properties specified in Dutton’s twelve item list. Artworks have aesthetic properties in virtue of having certain combinations of the properties listed in the definition.

Such a view is compatible with aesthetic functionalism, the theory that defines artworks as things endowed with aesthetic functions, whatever other functions they may have and whether or not their primary function is an aesthetic one. Both theories are compatible because aesthetic functionalism does not imply the view rejected by Dutton that artworks have aesthetic properties “by its nature” (essentially), only that they should or are supposed to sustain aesthetic properties (or to generate aesthetic experiences, depending on the particular version of aesthetic functionalism one adopts). As opposed to aesthetic essentialism, aesthetic functionalism can accommodate the existence of dysfunctional artworks.

An issue which is also relevant to the compatibility between Dutton’s cluster theory and aesthetic functionalism is whether the latter must yield a disjunctive definition in order to accommodate works that apparently lack an aesthetic point, such as ready-mades and much conceptual art. This will depend on how the aesthetic functionalist conceives of aesthetic properties, particularly, on how wide their dependence base is thought to be – does it include perceptual properties only, or does it also include contextual properties?

Construing Dutton’s list of criteria as a wide dependence base for aesthetic properties opens up what I believe to be a promising strategy for the aesthetic functionalist. This is connected with item 11 in Dutton’s list: contextual connections between a candidate artwork and extant art traditions and institutions. The strategy I’m envisaging explores the possibility that objects that don’t appear to have any aesthetic functions might acquire them when placed in an artworld context or having the status of artwork conferred upon them. In order for this strategy to avoid the charge of being ad hoc

⁶⁹ Dutton (2006: p. 374). The items in Dutton’s list are: 1) affording direct pleasure, 2) exhibiting skill or virtuosity, 3) style, 4) novelty and creativity, 5) being the object of criticism, 6) representation, 7) “special” focus, 8) displaying creative individuality, 9) emotional saturation, 10) posing an intellectual challenge, 11) existing in art traditions and institutions, 12) affording imaginative experience. *Ibid.* pp. 369-373.

there must be room for the possibility that some attempts at generating aesthetic properties this way may fail (otherwise it would be too easy to make aesthetic functionalism accommodate any purported cases of apparently non-aesthetic artworks). I will discuss this strategy in more detail when I present my reasons for rejecting a proceduralist theory. For now, the relevant point is that Dutton's cluster theory is compatible with such a version of aesthetic functionalism and is thus compatible with a more basic (aesthetic) definition of art that doesn't require formulation as a cluster definition.

The point just made about aesthetic functionalism and cluster definitions could be made for a proceduralist theory of art, as long as we accept, with Stephen Davies, that there is no inconsistency between a procedural definition of art and the view that art is importantly functional (the proceduralist need only deny that the functionality of art has any bearing on its definition). Given the primarily epistemological role of cluster accounts of art, and given that they are compatible with both essentialist and anti-essentialist underlying metaphysical views on art, I believe it's safe to assume that the real metaphysical game will be played out between simple and disjunctive definitions and anti-definitionalism.

2 Social ontology and art kind ontology

2.1 Function and procedure: causal-agentive art and status-functional art

John Searle (1995; 2010) has argued that institutional facts are distinguished from other social facts by the implicit application of a constitutive rule, according to which something (a physical object, a person, or a representation) counts as having the relevant status-function in a given context. The form of such a rule is the following: “X counts as Y in C”. Two important formal properties of this rule are the following: it iterates indefinitely upwards, and it interlocks indefinitely horizontally. What this means is that any Y term in a given application of the rule can be the X term in a further application of the rule (iteration). For instance, uttering a certain sentence (X) counts as making a promise (Y) in a given context (C); making a promise (X) counts as getting married (Y) in a certain context (C). In the case of marriage, the relevant context may include something like a ceremony being conducted by a priest. And being a priest consists in someone (X) counting as a priest (Y) in a given context (C). So the context variable in each application of the formula (the C term) will typically include other institutional entities. Any institutional fact one can point to will be connected to an indefinite number of further institutional facts (interlocking).

Take the simple case of a stone wall functioning as a territorial boundary and the line of stones resulting from the gradual decay of the wall over times: a stone wall (X) counts as a territorial boundary (Y) in a context (C). When the stone wall decays to the point that the causal-agentive artifact it once was no longer exists, the resulting line of stones still has the relevant status-function in virtue of being collectively represented as a territorial boundary. Institutional reality has the following interesting ontological property: something is an objective institutional fact only because people believe it to be so; because there is a convergence in how a certain aspect of the world is represented. When the shared beliefs in the system of status functions are radically questioned, the existence of the institutions is endangered. This explains an important point for the institutional theory of art: before the emergence of the twentieth century artworld, it is hardly conceivable that anyone could make it the case that a ready-made was an artwork. This is sometimes taken to be a point in favor of the procedural theorist. For something to be an artwork, the artworld must be ready for it. Different procedural theories account differently for this fact. It is either that some new art theory must be in place in order for the object to be accepted; or the institutional frame in which ready-mades can be artworks must be in place. To the proceduralist, it is not enough that John Cage presents his gurgling of carrot juice as a composition. When he does this, his intentions are already

adequately framed by the institutions of the artworld, allowing him to bring the conventions of the artworld to bear on his composition. The same intention of presenting that event as a composition could not be successful in earlier historical stages of the artworld.

We can now apply Searle's conceptual scheme to the institutional theory of art. Take the example of the most well-known of Duchamp's ready-mades, *Fountain*. The urinal (X) counts as an artwork (Y) in the twentieth-century artworld (C). In this scheme, artwork status is a special case of status-function artifactuality. Nothing is an artwork in virtue of its physical properties alone. A convergence in representation must obtain for something to perform the artwork role in a given context. Likewise, someone is not an artist in virtue of his physical or biological properties alone. An artist is someone who is recognized as such within a social practice (the relevant context). The totality of the interlocking status-functions will constitute what is called "the artworld". Now we have a clearer formulation of the institutional theory *and* a clearer notion of what makes an institutional theory institutional: its reliance on status-functions. Once again, what distinguishes the institutional theory from the aesthetic theory of art is not the appeal to functions, but the appeal to a special *kind* of functional properties. An object becomes an artwork by the same mechanism in virtue of which a sound/squiggle becomes a word in a language or in virtue of which uttering certain words in a context counts as getting married.

The appropriateness of this redescription of the institutional theory of art in terms of Searle's theory of institutional reality can be captured in the way George Dickie himself has replied to Richard Wollheim's (1980) well-known challenge to the institutional theory: the so-called "Wollheim's Dilemma". Either the representatives of the artworld have *good reasons* to bestow artwork status upon an object or they have no reasons to do so. If they have reasons, then it is those reasons that should figure in a definition of art; if they don't have any reasons, then the concept of *artwork* is an arbitrary concept. Dickie responded to this challenge in his paper, "Wollheim's Dilemma", by analyzing a parallel case of status conferral:

"Presumably having a good reason is required for a king or queen to confer a knighthood. Say that a man has had knighthood conferred on him because he is believed to have slain a dragon. Believing the man to be a dragon-slayer is the good reason that he has had knighthood conferred on him, but he would not be a knight if a king or queen had not conferred the status on him. The conferring does not 'drop out' as a necessary part of becoming a knight because it is done for a good reason. That conferring 'drops out' if a good reason is necessary is just false as a generalization."⁷⁰

⁷⁰ Dickie (1998: p. 131).

This response is adequate and is precisely what one would expect if artwork status is thought of in terms of a status-function. There will be many different reasons to ascribe any given status-function to an object, person or representation⁷¹ and there is no reason to suppose that the relevant reasons will make up a well-defined *kind*. But whether something *has* the relevant status is not a matter of what reasons were effective in motivating the agents conferring the status, but of the *procedure* by which the status was conferred upon the object. It is perfectly conceivable that there are good reasons *not* to draw territorial boundaries in such and such a way. This fact doesn't make it less the case that those boundaries *are* in fact boundaries. The existence of reasons to confer (or *not* to confer) art status upon an object only constitute an objection to the institutional theory of art *if* it is the case that artwork status is *not* to be conceived in terms of status-functions; but then Wollheim's challenge begs the question against the institutional theory. Reasons for and against status conferral will not do the trick; an argument against arthood as status-functional is required. If all the reasons that artists have had, in the total course of art's history, to present an object as a candidate for appreciation, plus the reasons artworld members had to accept those objects, formed a well-defined set or consistently pointed to specific features of those objects (e.g. their aesthetic properties), we would have the beginning of an argument. But chances that there will be one single kind of reason or even a plausible, well-delimited disjunction of reasons for art status conferral in the course of art's history are dim at best. Reasons for art status conferral will be as diverse as the valuable things we can find in artworks, and the set of things that can count as artistic values far exceeds any kinds of value or purpose in which (causal-agentive) functionalist theories of art have relied (representation, expression, aesthetic value). For instance, one of the reasons John Cage's 4'33'' is a valuable artwork is that it makes us think about the nature of music. Why would that be a *bad* reason for art status conferral (even if it falls short of being a good reason to count Cage's piece as *music*)? Wollheim's move is question-begging because there is no way of discriminating between *good* and *bad* reasons for art status conferral *independently* of a theory of art. Thus, Wollheim's dilemma only poses a challenge to the institutional theory at the cost of sneaking some version of functionalism in through the backdoor, so to speak.

I believe a better argument for (causal-agentive) functionalism about art lies somewhere else. But before getting into it, we must get rid of one assumption that has accompanied debates about

⁷¹ One interesting aspect of institutional reality is that in some cases, the X term does n't even need to be a *physical* item at all. A case in point is that of money: most money nowadays has no physical existence at all; it is no more than representations of a certain kind of power in the form of computer digits.

the nature of art and is shared by both “essentialists” (those who defend one or other definition of art either in terms of a conjunctive or a disjunctive definition) and “anti-essentialists” (those who argue that art can’t be given a definition in terms of necessary and sufficient conditions). This assumption is *art concept monism*; that is, the thesis that there is one, and only one *art* concept that can or cannot be defined in virtue of its structure. So, accordingly, anti-essentialists since Weitz argue that our single concept *art* has a certain structure – it is what they call a “family resemblance” concept – and essentialists of all colors, on the other hand, think that our single concept *art* can be defined, either in terms of necessary, conjunctively sufficient conditions, or in terms of disjunctive sufficient conditions, none of them necessary.

Christy Mag Uidhir and P. D. Magnus (2011) have recently argued for *art concept pluralism*, which is the view that there is no single *art* concept, but a plurality of concepts, each one of them more or less useful depending on the *purpose* it serves. They argue that we should think about the *art* concept in philosophy analogously to the way the concept *species* is treated in biology. They isolate three *species* concepts used by biologists: 1) the *phenetic species* concept we get from eighteenth century taxonomy (using this concept, we sort species in terms of their exhibited or observable features), 2) the *biological species* concept, introduced in the twentieth century (by which we sort species in terms of reproductively isolated groups), and 3) the *phylogenetic species* concept (which sorts species in terms of the smallest group of common descent). Each of these concepts presents its own difficulties for a unified view of species, but each is more or less well suited to different kinds of biological *enquiry*. I do not wish to enter into the discussion of the relative merits of each concept for biology, since that is a departure from my subject matter, but it is important to mention them as a case where concept pluralism is advantageously applied in a field of enquiry. Concept pluralism, as Mag Uidhir and Magnus argue, does not imply epistemological anarchy: “The fact that the word ‘species’ is used by some community to denote a concept does not automatically make that concept a legitimate *species* concept. For example, we can reject a naïvely essentialist *species* concept that rules out boundary cases and intermediate forms as conceptual impossibilities, a *species* concept that organizes trees by the type of lumber that may be wrung from them, or a *species* concept that organizes animals by how they taste.”⁷²

According with Mag Uidhir and Magnus’ proposal, there are at least four art concepts of distinct usefulness for different inquisitional or investigative purposes: 1) *aesthetic art*, 2) *conventional art*, 3) *historical art* and 4) *communicative art*. The first corresponds to artifacts that fulfill some

⁷² Mag Uidhir & Magnus (2011: p. 90).

aesthetic function and therefore satisfy different aesthetic definitions, in virtue of affording aesthetic values or experiences or of requiring an aesthetic attitude, and so forth. The second corresponds to artifacts that are recognized, accepted or governed by artworld conventions and institutions and therefore satisfy procedural definitions. The third corresponds to artifacts that emerge or are embedded in art-historical traditions or narratives and is useful for historical enquiries of the kind marshaled by Levinson (1990) and Carroll (1993). The fourth corresponds to artifacts that serve a communicative purpose (representational, semantic or expressive) and thus are captured by theories like that of Danto's later writings (1981) and John Dilworth (2004).

I want to focus on the first two concepts, since I view aesthetic functionalist theories and proceduralist theories as the two major competitors on the field. The reason I see these theories as the two major competitors is in part connected with Dickie's arguments for classifying theories of art into two major groups: NKTA and CKTA. In this classification, historical theories and Danto's later theory of art are hybrid theories. Furthermore, Danto's later theory is more obviously adequate to a specific kind of art (whereas it is not clear that the "difficult cases" for aesthetic functionalism and proceduralism cannot be adequately dealt within the scope of those theories), and Dilworth's theory has more important consequences for the ontology of art than for art's definition (he sees both singular artworks such as painting and non-cast sculpture, and multiple instance artworks, such as music and etchings, as things that are *represented*, not instanced, by their "vehicles" – in his view, the *Mona Lisa* is a proposition-like abstract object and the painted panel in the Louvre is a *representation* of that object). And since my ultimate goal is to provide a useful definition of *music*, not art (art is important to our enquiry in part because definitional and anti-definitional theories of art are the same that can be brought to bear on the nature of music) I leave historical theories more or less aside since even the major proponent of an historical theory of art, Jerrold Levinson, endorses a definition of music in aesthetic functionalist lines. Yet another reason to focus on the first two concepts is that historical theories can be seen more as theories about the *identification* (in the epistemic sense) of artifacts as artworks than as theories about the *metaphysical* status of arthood (the second major proponent of an historical theory, Noel Carroll, explicitly proposes his view in this way).

As I have pointed out, Stephen Davies, who endorses a version of the institutional theory, has diagnosed the current debate between what he calls "functionalist" theories of art and what he calls "proceduralist" theories, as having reached a point where both perspectives argue past each other, and regularly beg the question against each other "by taking for granted one thing or another which

the other side is committed to rejecting.”⁷³ We have already seen an example of this when looking at the so-called “Wollheim’s Dilemma”. The (causal-agentive) functionalist appeals to “good reasons” for status conferral, thereby assuming that all good reasons will be subsumed, say, in the class of aesthetic values; the proceduralist (or status-functionalist) has a good response, but nevertheless one that requires argument for the thesis that arthood is a status-function rather than a causal-agentive function, and has several problems dealing with first art and art which does not appear to be institutionally framed. I have argued that matters can be clarified by actually seeing both approaches as different kinds of *functionalist* theory, one relying on some causal-agentive function (such as aesthetic functions), the other relying on the notion of arthood as a status-function. I’ve motivated this view by providing a classification of functional properties and artifacts that are defined according to one or the other kind of functional property. This classification also fits well with Dickie’s scheme of NKTA and CKTA, with causal-agentive theories falling under the first group, and status-functional theories falling under the latter.

Other aspects of Davies’ diagnosis are the assumption, on the part of the (causal-agentive) functionalist, that a definition of art must reveal art’s purpose or ultimate value, and the further assumption that there is one single overarching purpose that art must serve. Accordingly, aesthetic functionalists about art often argue that the institutional theory fails for not providing an account of art’s value, dissociating art status from the value of art. So, what for the institutional theorist amounts to a theoretical virtue – that his theory provides an explanation of the classificatory sense of “art” as distinct from its evaluative sense, the aesthetic functionalist sees as a theoretical demerit: in his view, the institutional theory explains art status by providing an “unexplained explainer”, since by disconnecting art status and art value, we are left with no explanation of why art institutions emerged in the first place, and what distinguishes them from non-art institutions. By contrast, the institutionalist assumes that if an account of art’s value is embedded in a definition of art we will be left without resources to explain the existence of *bad art*, since he thinks that a consequence of an evaluative definition of art is that bad art is disqualified from the realm of art, for any object that fails to satisfy art’s purpose will not be classified as art. However, an evaluative definition of art is perfectly consistent with the existence of bad art, as Nick Zangwill (2007) pointed out, since there is a relevant difference between aesthetic *functionalism* and aesthetic *essentialism*. The envisaged consequence of an evaluative definition of art would only follow if artworks are conceived as having aesthetic properties essentially. But aesthetic functionalism is a

⁷³ S. Davies (1990: p. 99).

normative theory, which means that aesthetically dysfunctional artifacts will still be classified as art, because they are the kind of thing that *should* perform an aesthetic function; it does not follow from this that the object *actually* performs that function.

Going back to Davies' diagnosis, on the part of the proceduralist, we also have the assumption that if there is no one single overarching purpose that art must serve, and if art is to be defined procedurally, then we must deny that art is importantly functional (in the causal-agentive sense). This begs the question, according to Davies, since art's being importantly functional is consistent with the absence of a single pervasive need that art satisfies.

"Because procedure and function have drifted apart and because art is to be defined procedurally, not all artworks will be appropriately functional. This fact is not at odds with a functional characterization of the importance of art in general. Indeed, our reasons for adopting and operating the procedures in terms of which art status is said by the proceduralist to be conferred are rendered inexplicable without some acknowledgment of the functionality of art in general as giving cause for us to institute and maintain the relevant procedures."⁷⁴

While I don't endorse institutionalism as *the* overarching theory of art, I believe Stephen Davies points in the right direction: once the first objects properly termed "artworks" emerged, function and procedure (or in my terms: causal-agentive artifactuality and status-functional artifactuality) were liable to drift apart, and this is what in fact happened in the twentieth century artworld:

"artists and others have attempted to confer art status on "found" objects (such as pieces of driftwood) and on "ready-mades" (such as snow shovels), and in some such cases they did so not because they recognize in those objects aesthetically meritorious properties but rather, precisely (and perhaps perversely), because such items seemed to lack any aesthetically rewarding properties at all."⁷⁵

But Davies acknowledges that to adequately account for first art, we must conceive it in terms of what Dickie called the "natural-kind activity" of seeking experiences of the aesthetic qualities of objects and modifying existing objects in order to exemplify these kinds of properties. An object has an aesthetic function if among the intentions with which it was made figured an intention that it should display a certain quality or set of qualities, appropriately termed "aesthetic" (even if the

⁷⁴ S. Davies (1990: p. 104).

⁷⁵ S. Davies (1990: p. 100).

makers of such artifacts do not possess an explicit concept of the aesthetic, the same way that the painters at Lascaux didn't have an explicit concept of depiction, although they produced aesthetically arresting depictions on their cave walls). This view doesn't require that aesthetic functions be isolated from other functions, or that the primary function of the object must be an aesthetic one, or that it is made solely for "contemplation". All that is required is that the objects "betray" a concern for aesthetic qualities, besides other functions they may serve. Furthermore, some aesthetic qualities may be dependent on one or more of the objects' non-aesthetic functions (some qualities may be good-making features in one kind of artifact but not another). This is often the case with representational works of art, where many of its aesthetic qualities depend on the object having the representational properties it has.

What all this means is that a causal-agentive concept of *art* in terms of aesthetic functions is required to explain the emergence of art behaviors and art institutions, but that once an art practice is in place, the purposes with which art is made can diverge from the causal-agentive functions that account for art's inception. In Davies' words, "the human creation and appreciation of aesthetic qualities might be *historically essential* to the concept we have, though mention of such properties is not required in a philosophically acceptable definition of what makes any particular artwork (created after first art) what it is."⁷⁶

I have my doubts that the aesthetic has been in any way superseded by other concerns in the making of artworks. That might be true for an infinitesimal part of art creation in the twentieth century artworld, and even with that qualification it doesn't account for *all* "artworld art" in the contemporary western world. The view that the aesthetic in art has become historically obsolete (not that I am attributing that view to Davies) is inadequate when we consider the whole range of art production both within the boundaries of the "western" artworld and outside those boundaries, either in "western" or "non-western" art practices. Thus, Denis Dutton correctly (I think) observes:

"Aesthetics at the outset of the twenty-first century finds itself in a paradoxical, not to say bizarre, situation. On the one hand, scholars and aesthetes have accessible to them—in libraries, in museums, on the Internet, first-hand via travel—a wider perspective on artistic creation across cultures and through history than ever before. We can study and enjoy sculptures and paintings from the Paleolithic, music from everywhere, folk and ritual arts from over the globe, literatures, arts of every nation, past and present. Against this vast availability, how odd that philosophical speculation about art has been inclined toward endless analysis of an infinitesimally small class of cases, prominently

⁷⁶ S. Davies (1997: p. 33).

featuring Duchamp's ready-mades or boundary-testing objects such as Sherri Levine's appropriated photographs and John Cage's 4'33''.⁷⁷

By quoting this passage from Dutton I'm not suggesting, of course, that Stephen Davies' approach to art is narrow-minded. That would simply ignore his wide-ranging work on both western and non-western art and music, as well as his recent book on the relationship between art, the aesthetic and evolutionary accounts of both, which deals extensively with the topic of first art. But it is worth noting that the importance of a minute class of avant-garde experiments in more recent times (minute in face of the vastness of past and present art production that does *not* drift away from art's causal-agentive origin) can be easily overblown by the philosopher's attraction for polemic cases. Even granting that Danto's "indiscernible duplicates" are of crucial importance to the philosophy of art, in the sense that they make explicit important features that were previously neglected (and I have argued that this is precisely the case, though not with the results that the proceduralist wishes), it is a mistake to infer from them a conclusion so strong as that the aesthetic in art has become historically obsolete (once again, not that Davies endorses such a view, but his words *can* be interpreted in that way).

This point connects with the observation that artists pick and choose the philosophical doctrines that appear to serve their purely artistic purposes (when often the purposes will be compatible with the falsity of the doctrines). Think of Varèse's redescription of himself as "a worker in rhythms, frequencies and intensities" and of his art as "organized sound". It is not the case that a definition of music narrower than "organized sound" will exclude the kind of aesthetic experimentation with sound that Varèse engaged in. In this line, the purported renouncement of beauty and the aesthetic in the twentieth century is a case in point. To what extent is "anti-aesthetic" art *really* anti-aesthetic, rather than constituting the endorsement of one set of aesthetic values by contrast to other sets of aesthetic values? To what extent isn't the "rejection of beauty" more like a rejection of "prettiness" (which is a kind of "substantive" aesthetic property) and less an actual rejection of aesthetic value in general? (Marinetti's "aesthetics of the machine" as well as Soviet "constructivism" come to mind here.) This seems to fit well with most historical cases of "artistic revolutions", in which departure from tradition is not a matter of simply leaving the aesthetic aside but of widening the scope of our experience, enabling new ways of experiencing the world *aesthetically*. In this sense, artistic revolutions are not radically different from the experience of immersion in the art of a foreign culture, a consequence of which may be our attunement to kinds of beauty to which we were

⁷⁷ Dutton (2006: p. 368).

hitherto oblivious or unable to experience (I think here of a western listener's evolving experience with, say, Gagaku or Gamelan music). The same "reshaping" of the sensibility was effected by the emergence of jazz music, which transformed our experience of harmony and exerted a powerful influence in "classical" composers such as Ravel and Stravinsky, who quite literally taught us new ways of listening to music. The vast majority of "artistic revolutions", even those that explicitly part with some aspect or aspects of the tradition in which they arise, are better described in terms of a broadening of human experience *by way of* aesthetic experimentation. This does not amount to a rejection of the aesthetic qualities in terms of which former art was appreciated (e.g. beauty, elegance, delicacy, balance, boldness, powerfulness, vividness, unity, coherence, gracefulness, etc.) and the way clusters of these qualities combine and interrelate to form a consistent whole, but to a realization of how those very same qualities can be instanced in an immense variety of ways (thus the broadening of experience: becoming able to experience forms of beauty to which one was hitherto oblivious plausibly impacts how one perceives and understands the world to as great an extent as the widening of the language resources with which one describes the world). One should also take care not to confuse the "eschewing of the aesthetic" with an exploration of how combinations of *negative* aesthetic qualities (ugliness, deformity, bleakness, gaudiness, garishness, etc.) can produce works of great aesthetic interest or value. This is an example of how "eschewing beauty" (or perhaps more appropriately, "eschewing prettiness") does not amount to a rejection of the aesthetic – there is more to beauty than pretty or "refined" things. In fact, it is only in the twentieth century western artworld that the issue of "anti-aesthetic" art clearly emerges as a problem for the philosophy of art. But such issues in turn could only arise because artists began to entertain higher order thoughts about their own aesthetic insights (which, as I have stressed here and there, they need not conceptualize in terms of an explicit concept of the *aesthetic*) and proceeded to create works that pose a challenge to our shared conceptions of art and the aesthetic – which is partly the intended function of ready-mades and other pieces of conceptual art, such as Cage's 4'33'', whose value consists partly in making us think about the nature of music and the aesthetic qualities of sounds we resist to categorize as music. At this point I would also venture that the "rejection of beauty" or the "rejection of the aesthetic", when explicitly endorsed by artists, may turn out to be confused expressions of more ideologically motivated programs, such as a rejection of "bourgeois culture", which are equally well represented in opposite regions of the ideological "spectrum": think once again of (precisely) the *aesthetics* of twentieth century futurism and Russian constructivism. Thus, what in fact is often the developing of new aesthetic sensibilities is easily confused with a "rejection of the aesthetic". In this line of reasoning, Nick Zangwill has observed that:

“if we look carefully even at the most radical ‘performance art’, ‘conceptual art’ and the like, we find that there *is* in fact some aesthetic concern. We should not believe everything that artists say about their work. It is safer to look at what artists *do*. We can ask: if conceptual art is so conceptual, why do we need to be provided with anything to look at at all? Why can’t it all be done with concepts? The truth, it might be said, is that the art is not entirely conceptual.”⁷⁸

About “anti-aesthetic” intentions in the production of art, Frank Sibley, in a paper about the relationship between the concepts of art and the aesthetic, has remarked that so-called “anti-aesthetic art” is, in a way, parasitic on more basic conceptions of art in terms of the aesthetic:

“Attempts to produce works that lack aesthetic interest are parasitic on the traditional concept. For it is as much a logical impossibility to intend something that lacks aesthetic interest without having a concept of the aesthetic, as to intend something that has aesthetic interest. Equally, the intention not to produce art, or to produce what is not art, requires a conception of art. So, it seems, whether we allow these modern works to be new forms of art or not, they all presuppose a familiar concept of art and therefore of the aesthetic.”⁷⁹

This is a special case of the more general point that an intention *not to do X* is still an intention defined in terms of a concept of *X*. So, even granting that artworks made with “anti-aesthetic” intentions are successful in actually not possessing any aesthetic profile (as opposed to being aesthetically good, aesthetically bad or aesthetically bland) their value, whatever it is, will still be crucially connected with the aesthetic, to the extent that it requires us to *think* about them in terms of the aesthetic, and so whatever value they have is parasitic on aesthetic value. It therefore seems that a consequence of this is that genuine challenges or counterexamples to an aesthetic theory of art would have to come not from “anti-aesthetic” artworks but from artworks whose causal histories trace back to no aesthetic intentions whatsoever, be they *positive* or *negative* aesthetic intentions.

Anyway, we don’t need to settle the score definitively in favor of aesthetic theories of art, for this is precisely what the strategy from art concept pluralism makes unnecessary. Whether or not the phenomenon of divergence between “function” and “procedure” (i.e., between causal-agentive function and status-function) is as narrow and parochial as the aesthetic theorist would like to show,

⁷⁸ Zangwill (2007: p. 65).

⁷⁹ Sibley (2001: p. 138).

whether or not “anti-aesthetic” artworks can be integrated by an aesthetic theory of art, Davies is right to observe that “procedure” doesn’t *have* to follow “function”. Rather than exclude or “colonize” the “hard cases”, the aesthetic theorist should enhance his theory to account for the possibility of divergence while maintaining the metaphysical and explanatory primacy of the aesthetic concept of art.

Searle’s theory of institutional reality provides us with a tool to accommodate this drifting apart of “function” and “procedure” in art practices. A crucial feature of that theory is the role of *language* in the creation of institutional reality. According to Searle, a mistake in which many philosophies of society have incurred is that of postulating the existence of language speaking individuals that then set about to create the first institutions and thus the fabric of social life. The reason this is a mistake is that if individuals share a language, they already have “social contract”; that is, they already have institutional reality, for language is the primary form of social institution. In other words, language itself is already a system of status-functions: words and sentences are status-functional artifacts; they do not discharge their representational functions merely in virtue of the physical properties of word or sentence tokens, but in virtue of shared representations. Language differs from other social institutions in two relevant aspects: 1) there was no point in time where people convened to decide upon the conventions for language, since this thought is viciously circular: in order to convene and make collective decisions, they would already have to share a language; 2) without a shared language there can be no other social institutions, since any application of the “constitutive rule” for creating institutional facts – “X counts as Y in C” – presupposes the representational abilities afforded by a language with the properties of discreteness, compositionality and generative syntax. Although people don’t have to explicitly think in terms of the constitutive rule in order to create institutional reality, they would not be able to share the required representations for it to be the case that a certain X counts as Y in a given context C, if they lacked a language. For it to be the case that something “counts as” having a certain property that it cannot have purely in virtue of its physical constitution, it has to be collectively represented in that way, and to represent something in that way language is required. An intuitive way to put this is to think about the following: we can imagine a society in which there is language but no marriages, no money, and no private property, though we can’t imagine a society in which there are all these things but not language.

Words differ from the vast majority of other status-functional artifacts in that they are “unintended emergent artifacts” (they also differ in this aspect from all causal agentive artifacts,

whose functions are intended by their makers). There are also unintended emergent artifacts that are not status-functional, e.g., hill paths and villages⁸⁰ – things that come to exist by a convergence in people’s behavior, without any of the participants ever intending the end result (they merely intend to walk across hills or build houses in a certain place). Most of the time, conventions arise not because anyone intended them, but because a certain “stable point” is spontaneously reached as a solution to a coordination problem. Take the example of the handshake:⁸¹ we learn how to use the handshake in a similar fashion to how we learn the use of words – by immersion into a community of language speakers. The handshake is an institutional artifact (it performs functions beyond the physical properties of its tokens: it can constitute a speech act or part of a speech act); it is also “abstract” in the sense that it recurs in space and time (it has multiple instances). The handshake convention most likely began in the same way many words come into existence: as an unplanned convergence in behavior. While some words are stipulated and explicitly introduced into a language, it is impossible for all words to have been generated in this way. Words are artifacts of a conventional sort and they obey the constitutive rule for institutional facts: a sound/squiggle (X) counts as a token of word W (Y) in the context of language L (C). But there was no point in time where a group of non-linguistic human beings convened to institute a language. Another important aspect of institutional facts that distinguishes them from other social facts is that their point is to generate positive and negative powers (rights and obligations). A person making a promise places herself under an obligation, that is, she creates a reason for acting which is independent of the actual set of desires and interests she previously had or already has. That some sounds and squiggles count as words and sentences generates a power of representing states of affairs, much beyond what any cognitively sophisticated animal is able to do without language. The point of creating institutional facts is to endow objects, people, events, etc. with properties and powers they can’t have merely in virtue of their previous set of intrinsic properties. And doing this is valuable because, like having language, it enables us to solve coordination problems.

Another crucial point in Searle’s theory of institutional reality is that once we have language, the primary social institution, the creation of further institutional facts is spontaneous and unavoidable: “Given a language you can, so to speak, create institutional facts at will (that is the top-down part); but when you have a language, other social institutions will inevitably grow up out of language (this is the bottom-up part).”⁸² This connects to what was said in 1.2 about culturally modulated natural-

⁸⁰ The examples, and the terminology, occur in Zangwill (2015: p. 135).

⁸¹ This example also occurs in Zangwill (2015: *ibid*).

⁸² Searle (2010: p. 63).

kind activities: language speaking beings such as ourselves will inevitably add layers of institutional reality to practically all our natural-kind activities. There is an important distinction, in Dickie's classification of theories of art, between cultural activities that are to be ultimately explained in terms of some NKA (e.g. the case of depiction and *seeing-in*) and cultural activities that are ultimately to be explained in the terms of a CKTA. This already allows us to see more clearly a deficiency in Dickie's treatment of the subject, a deficiency that "contaminates" his institutional approach to art and lead him to abandon, in his later formulation of the theory in 1984, the clause about *status conferral*. This move was rightly criticized by Stephen Davies (1988) as amounting to an abandonment of the institutional approach in favor of a mere "social context" approach. That art requires a *social context* does not by itself show that art is institutional. An institutional theory requires a clear distinction in social *roles* and the exercise of authority (deontic power) by those who are in a position to bring the conventions of art to bear on "candidates for appreciation". It requires a notion of how status-functional artifactuality is created, in terms of this act of creation being itself embedded in a system of status-functions. Dickie abandoned the status conferral clause on his early formulation of the institutional theory after a host of criticisms to the effect that the "artworld" does not act as a single individual and that no formal ceremonies of status conferral are ever enacted to endow an object with artwork-status. In part due to such criticisms, Dickie eventually moved to a formulation of the theory according to which "an artwork is an artifact of a kind to be presented to an artworld public". This definition, though, is compatible with a non-institutional interpretation, since not everything which is part of culture is thereby institutional in nature. Davies claims that Dickie should have kept the status conferral clause. I agree. The early criticisms of the institutional theory were misguided. They result from an incomprehension of the nature of institutional reality and thus show how social ontology is required if we are to make progress in the metaphysics of art kinds. In other words, it vindicates the metaphysical constraint: "No art ontology without social ontology".

2.2 Institutions, games and status-functions

Arthur Danto's 1964 paper titled "The Artworld" inaugurated what has come to be known in the literature as *proceduralist* theories of art. As we have already seen, what marks a theory as proceduralist is its contention that an artwork derives its art status from its being conferred upon it in accordance with a *procedure* and not from the work's effectiveness in fulfilling (or its being supposed to fulfill) a goal (or set of goals) specific to art. Proceduralism is traditionally opposed to functionalism, for which art status derives precisely from such an art-specific goal.

In the above mentioned paper, Danto made his influential claim that what makes artworks *artworks* is not a feature that would perceptually distinguish a non-art object from an art-object, something like a multiply realizable "significant form", as was postulated by Clive Bell (a kind of visual or aural configuration that has the effect of "drawing us toward" the object and infuses us with a particular emotion or pleasure), but "an atmosphere of artistic theory":

"To see something as art requires something the eye cannot decry – an atmosphere of artistic theory, a knowledge of the history of art: an artworld."⁸³

Phrased in this way, the thesis is not as strong as it may seem, for one could reply that even though to *see something as art* (or to conceptualize it as art) requires a context of art theory, seeing something as art is not what makes something art. Another difficulty is this: if seeing something as art requires art theory, where did our art theories come from? (After all, they are theories *about art*!) I'll ignore this difficulty for now. A stronger formulation of the same view comes up further ahead in Danto's paper:

"It is the role of artistic theories, these days as always, to make the artworld, and art, possible. It would, I should think, never have occurred to the painters of Lascaux that they were producing *art* on those walls. Not unless there were Neolithic aestheticians."⁸⁴

Once again, one could argue that the inability of the cave painters to see (or conceptualize) the result of their efforts as art is compatible with those things, nevertheless, being art. After all, the stone implements that those ancestors of ours used in their daily tasks, as well as the tools they used

⁸³ Danto (1964: p. 580).

⁸⁴ *Ibid.* p. 581.

to make cave paintings, were already *artifacts* even before human beings conceptualized them as artifacts (where “conceptualizing them as artifacts” is distinct from having the intentions requisite to make an artifact, e.g. the intention to make a cutting tool, and the beliefs and desires which in turn are required to have those intentions). An argument is needed to the effect that the category *artwork* is relevantly different from the category *artifact*, in that the production of artworks, unlike the production of non-artistic artifacts, requires a *concept* of artwork. Otherwise, one can even contend that “seeing something as art” doesn’t require explicit artwork concepts; all that is needed, it may be claimed, is that the cave painters regard their paintings as something distinct from their other artifacts, where a disjunction of possible “acts of regarding” that don’t necessarily deploy an explicit artwork concept would be sufficient. This would mirror the situation where in order to regard an object as an artifact it would suffice to see it as something one has made to facilitate a certain task.

According to the stronger formulation above, artistic theories donot just enable objects to be seen as artworks; they are what makes art possible. The objections raised above can be seen as a charge of confusion between the epistemological and the metaphysical levels of description, so to speak. But a Danto-style proceduralist could simply respond by biting the bullet and claiming that in the case of artworks, epistemology and metaphysics are inextricably bound. For there to be artworks just is for certain objects to be seen as artworks, and this requires a background of theory, a context.

Danto went on to develop and defend his idea through a series of thought experiments that originated in actual encounters with real life artworks – Warhol’s *Brillo Boxes* and Duchamp’s *ready-mades* – though perhaps the most widely influential is that of a fictional work, imagined in Jorge Luis Borges’ short story about Pierre Menard, author of *Quixote*: a fictional tale of a nineteenth century text which is word per word identical with the text of Cervantes’ work, *Don Quixote*.

The key concept in all these thought experiments is that of *indiscernible duplicates*: pairs of perceptually indistinguishable objects which, in virtue of differing contextual properties, either constitute distinct artworks sharing identical *vehicles* (the indiscernible objects in which both artworks are “realized”) or a case where, again in virtue of differing contextual properties, an artwork is visually or aurally indistinguishable from a non-art object (what Danto calls “a mere real thing”). Accordingly, the fictional text of Pierre Menard and Cervantes’ novel constitute two distinct artworks that share the same vehicle (a certain sequence of words). They are distinct artworks because their originating contexts endow them with distinct, incompatible properties. On the other hand, Duchamp’s *Fountain*, an “appropriated” urinal which the artist submitted anonymously to the 1917 exhibition of the Society of Independent Artists, an artwork, shares its

physical structure and perceivable properties with a vast array of non-art objects (common urinals) that were never “appropriated” into an artworld context.

In his 1981 book, *The Transfiguration of the Commonplace*, Danto presented a series of thought experiment cases, such as that consisting of eight perceptually indiscernible red canvases of square shape:⁸⁵ One of the canvases is titled *Israelites Crossing the Red Sea*, and it represents the moment after the Israelites had crossed over and the Egyptian army had drowned; another canvas is titled *Kierkegaard's Mood*; a third canvas is titled *Red Square*, representing “a clever bit of Moscow landscape”; a fourth canvas, though sharing the title *Red Square*, is “a minimalist exemplar of geometrical art”; a fifth canvas is titled *Nirvana*: “It is a metaphysical painting based on the artist’s knowledge that the Nirvanic and Samsara orders are identical, and that the Samsara is fondly called the Red Dust by its deprecators.” A sixth canvas is titled *Red Table Cloth*, and consists of a still-life by “an embittered disciple of Matisse”; the seventh canvas is “not really an artwork, but merely a canvas grounded in red lead upon which, had he lived to execute it, Giorgione would have painted his unrealized masterwork *Conversazione Sacra*”. The final canvas is neither an artwork nor has it art-historical interest; it is “a surface painted, though not *grounded*, in red lead: a mere artifact.”

This imaginary exhibition contains eight objects which are facsimiles. Six of them are distinct artworks; two are distinct (non-artwork) artifacts, one of which has some art-historical interest while the other has none. In four of the artworks (canvases 1, 3, 5 and 6) the red color has a *representational* function; in one of these cases – *Nirvana* – the representation is metaphorical; in the canvas 2 – *Kierkegaard's Mood* – the color has an expressive function;⁸⁶ in the minimalist painting redness is referred to by exemplification (in Goodman’s sense).

It was from observations of this kind that Danto inferred his theory of art as *interpreted* objects (“mere real things” or “vehicles”), and his ontology of artworks as consisting of two-part entities: a “mere real thing” and an *interpretation*. Danto thought that the existence of *art* depended on “theories of art” since, as he saw it, the difference between artworks and “mere real things” (or “mere real things” which are also artifacts) is that artworks have the essential property of “aboutness”, a property that they have partly in virtue of standing in a relation to a “theory of art”. Ready-mades, for instance, could not have been artworks in previous “stages” of development of the artworld, since only in the twentieth century artworld could they be seen as having this essential property.

⁸⁵ Danto (1981: pp. 1-2).

⁸⁶ See D. Davies (2004: pp. 63-64).

Danto's theory is thus an example of non-institutional proceduralism. The existence of artworks is indexed to the existence of theories of art, and these depend on social practices, but the relational properties that are of definitional and ontological importance to artworks are not relations to particular institutions or institutional acts of status-conferral.

Danto's particular version of proceduralism differs from another well-known representative of the class: the institutional theory of art, as was originally developed by George Dickie in two distinct versions, over an interval of ten years, between 1974 and 1984.⁸⁷ According to the first version of the institutional theory "a work of art in the classificatory sense is 1) an artifact, 2) a set of aspects of which has had conferred upon it the status of candidate for appreciation by some person or persons acting on behalf of a certain social institution (the artworld)"⁸⁸

The later version will drop reference to the conferral of status of candidate for appreciation, defining artwork as an artifact, made by an artist, which is of a kind to be presented to an artworld public.⁸⁹ This version of the theory is supplemented with a set of interdependent definitions of "artist", "artworld public", "artworld" and "artworld system", so that an artist is "a person who participates with understanding in the making of an artwork", an artworld public is "a set of persons the members of which are prepared to some degree to understand an object which is presented to them", the artworld is "the totality of artworld systems", and an artworld system is "a framework for the presentation of a work of art by an artist to an artworld public."⁹⁰

Many criticisms have been raised against both versions of the theory, but I'm not interested in going through all of them here. What interests me is, on the one hand, to isolate the feature or features in virtue of which a theory of art is not only a *proceduralist* one but specifically an *institutional* theory. The resultant objection will not be an objection specifically directed at Dickie's or any other particular institutional theory but rather an objection to the very idea of defining art institutionally. This objection, I believe, will be a particularly strong one, given that it integrates a criticism of Dickie's version of institutionalism by a friend of the institutional theory, Stephen Davies, namely, that a theory of art, if it is to be an *institutional* theory and not simply a social-context theory,⁹¹ must preserve the idea that art status is *conferred* upon an object by someone with the appropriate credentials. The idea of status conferral was dropped in Dickie's second version of his theory, in face of criticisms to the effect that the artworld does not act as a single individual,

⁸⁷ Dickie (1974; 1984)

⁸⁸ Dickie (1974: p. 34).

⁸⁹ Dickie (1984: p. 80).

⁹⁰ *Ibid.* pp. 80-82.

⁹¹ S. Davies (1988: p. 315): "The institutional theory must hold that the Artworld is a structured institution (and not merely an "institution" to the extent that it is social)."

investing objects with artwork status in something like a formal ceremony. As it will be made clear, on the light of a characterization of institutional facts, these criticisms are wide of the mark. The institutional theory has resources to escape the objection while preserving the notion of status conferral, but it will face a more serious objection.

The endeavor will not be a purely negative one, though, since there are crucial intuitions on the side of the proceduralist that must be assimilated by a functionalist theory of art if it is to make a significant progress in the current debate. Three of these intuitions are the following: 1) that artworks have important conventional aspects and a purely functionalist theory will not be adequate unless it can handle these aspects in an appropriate fashion, instead of dismissing them as idiosyncratic quirks; 2) that many things that are accepted as artworks in one social and institutional context would not be intelligible as such in all contexts and historical periods; 3) that even if a functionalist concept of art is required in order to account for how art practices get started (otherwise art theories and artworld institutions will turn out to be “unexplained explainers”)⁹², given that function and procedure can always become at odds with each other (artists may force them to come apart, e.g., by intentionally producing works that eschew any aesthetic point), an enhanced functionalist theory must give a satisfactory explanation of how this is so.

Stephen Davies has argued that the dispute between proceduralism and functionalism has fallen into a sort of stalemate, where both theories simply argue past each other. As Davies sees it “each [theory] regularly begs the question against the other by taking for granted one thing or another which the other side is committed to rejecting. In this way both sides regularly fail to engage with the concerns of their opponents, and the polemical dust raised by their clashes obscures the bases of their disagreements.”⁹³

We have seen that there are two major categories of artifacts, in virtue of there being two major kinds of artifactual functions that particular artifacts can perform. These two categories are a) *causal-agentive* and b) *status-functional*. They are distinguished by the kinds of relational properties an object must have in order to fall under one or the other artifactual category. Specifically, status-functional artifacts depend, not only causally but conceptually, on convergences in behavior (most importantly, convergences in ways of representing reality). My main example of a status-functional artifact has been that of *words*. No physical object or event (e.g. squiggles or sounds) counts as the tokening of a word unless there is a *shared representation* generating the

⁹² Zangwill (2009: p. 39): “Either the social practice is explained or not. If not, the existence of the institution is an unexplained explainer. This is unsatisfactory. The existence of the institution needs to be explained.”

⁹³ S. Davies (1990: p. 99).

norms according to which a certain pattern of physical marks or sounds will be recognized as a token of that particular word. That is why we say, e.g., that “duck” is a word in English but “klurg” is not. I can’t make it the case that “klurg” is a word in English just by using those marks or the correspondent sound pattern with the intention to refer to something. There would have to be a convergence of speakers in the use of that pattern so that “klurg” could function as a word, as a vehicle of meaning. Of course, I could introduce the concept KLURG by stipulating a definition of it (many concepts are introduced in this way), but that would only be effective if the use was taken up by other speakers. Plus, I’m only able to stipulate a meaning for “klurg” in virtue of my ability to use many other words that were not introduced by stipulation. One needs (public) words in order to make such stipulations. If I use “klurg” privately, with a personal meaning unknown to everyone else, “klurg” is not functioning as a word at all, since it does not enable me to make my thoughts known to others. Something functions as a word only when it constitutes a solution to a coordination problem. In fact, thinking about “private words” as *words* is only intelligible given the fact of public language. In this sense, my private “words”, like “klurg”, are a different kind of artifact altogether. Maybe it can function as a sort of mnemonic device. Words (in the primary sense) can serve that function, but the function of being a mnemonic device is distinct from the functional properties in virtue of which words are words. The idea of words before public language is like the idea of there being some unruly “chess” movements before the constitutive rules of chess are invented. Private “words” are no more words than Monopoly “money” is money; in fact the latter is closer to money than private “words” are to words, since at least it has the same kind of artifactual function: it is recognized by players of the game as a means of exchange; “real” money differs from make-believe money only in the extent of the make-believing that takes place, or in the spatiotemporal boundaries of the game.⁹⁴ The idea of private words, thus, shares the same property that gives the idea of “private money” a smack of weirdness: the only reason we’re able to think of some random item as “private money” is our previous grasp of the (status)functionality of money properly so named. Since being money *consists* in being a (special sort of) means of exchange, the concept of “private money” is pointless, because such thing cannot *mediate* anything at all (the buying and selling of goods and services is not a relation in which one engages with oneself). What

⁹⁴ Although it’s not accurate to say that we “make-believe” that some discs of metal are real coins, there is something akin to make-believe in “real” cases of metal discs or pieces of paper that are money: they are only money because we believe they are. It is the same kind of game, but with different boundaries. In fact, games are prime examples of institutions: things that only exist in virtue of constitutive conventions. A metal disc *counts* as money in much the same way that a certain piece of carved wood *counts* as a pawn or a knight and a certain movement of carved wood pieces on a chequered board counts as a legal move in chess: through a convergence in representing those things as such.

makes an object count as *money* (whatever concept we use to refer to it) is its occupying a certain place in certain relations between agents. *Mutatis mutandis*, the same is true of words.

Money is thus another prime example of status-functional artifact. It is an objective fact that some discs of metal (or pieces of paper, or wampum shells, or whatever) count as money. And what it means for a disc of metal to count as money is that people *converge in representing it* as money; shared beliefs are the truthmakers of sentences of the form “X is money”. Objects that count as money are not selected to perform that function in virtue of their pre-conventional efficacy in serving that purpose (or a relevantly similar causal ancestor of that purpose). The purpose (the proper function) in this case is generated by the convention, i.e., by the convergence in behavior; it only exists within a certain network of relations and it is not even intelligible if we conceptually disconnect it from that network of relations.

The crucial concept in delineating the category of status-functional artifacts is that of a *status-function*: a function that depends on a collectively recognized status. This was, to my knowledge, first articulated by John Searle in his work on social ontology. The concept of a status-function contrasts with the concept of a *causal-agentive function*: a function that depends not on a recognized status but on a causal relation obtaining between an object with an appropriate physical structure and the intentions and insights of an agent who puts that structure to a use (gives it a function), whether or not he actually produced the structure by shaping some pre-existent material. The appropriateness of the structure depends on the intended function, and the fit between them is also a standard by which we judge the appropriateness of the intentions. The relevant intentions are intentions of acting according to an insight that certain physical properties are adequate to implement a certain function. The content of the insight is a dependence relation between functional properties and physical properties. If someone intends to create a coffee-maker and the result is something physically indiscernible from a screwdriver (the example is bizarre but illustrative), something must have gone wrong, either at the level of insight (the envisaged physical properties do not sustain the envisaged functional properties) or of execution (the agent failed in producing the envisaged physical properties).

This picture of how a causal-agentive artifact comes into being parallels Nick Zangwill’s aesthetic-functionalist theory of artworks, which, according to functionalist theories generally, are a subclass of causal-agentive artifacts: an agent has an insight that certain non-aesthetic properties will sustain certain aesthetic properties (in other words, that a *dependence relation* obtains between the latter and the former) and sets out to implement those aesthetic properties by producing or appropriating an object with the relevant non-aesthetic properties, as envisaged in the insight. Different (causal-agentive) functionalist theories of art will appeal to other kinds of causal-agentive

function (e.g. Tolstoy's and Collingwood's expression theories of art). Aesthetic functionalism is distinguished from other kinds of functionalism about art by its specific appeal to aesthetic functions, which are rooted in aesthetic insights and intentions, which are intentions to act on those insights in the production of an aesthetically functional thing: an artwork.

Aesthetic functionalism, thus conceived, is a special case of causal-agentive functionalism, which subsumes all causal-agentive artifacts, artworks and non-artworks. If an artifact has at least one aesthetic function, it will fall under the category of *artwork*, within the wider category of *causal-agentive artifact*. The two major kinds of functional properties – causal-agentive and status-functional – determine the two major categories of artifacts, namely, causal-agentive artifacts and status-functional artifacts. This yields two different explanations of the nature of artifacts, according to the kind of functional properties they exemplify: a) causal-agentive functionalism; and b) status-function functionalism (if I may be allowed the inelegant expression). Partial overlapping of both categories of artifacts is also to be expected, which yields cases of artifacts with both kinds of functions, where the functions are independent of each other, and cases where one kind of function is dependent on the other. This means that the members of a given sub-category of artifact (e.g. *artwork*) can belong to *both* major categories simultaneously. For instance, (assuming the truth of aesthetic functionalism) an artifact which is also an artwork may be an artwork *in virtue of* having certain causal-agentive functions (namely, its aesthetic functions); but the very same artifact can, at the same time, have a status-function which is connected with the causal-agentive function in virtue of which that artifact is an artwork. In other words, the artifact *counts as* an artwork (it is recognized as such) in virtue of satisfying the causal-agentive concept of art (i.e. in virtue of having at least one aesthetic function). This is the case for societies that recognize the category of artworks as a distinct category (i.e. societies that have a concept of art), but in which “function and procedure” have not “become at odds with each other” (e.g. the western artworld before the twentieth century). It must be kept in mind that artifacts can satisfy the causal-agentive concept of art even in societies that have no concept of art. In such cases, it is *we* who count such artifacts as artworks. The phenomenon which Stephen Davies has described as function and procedure “becoming at odds with each other” can be described using the framework I've developed in the following way: an artifact *counts as* an artwork, not in virtue of satisfying a (more basic) causal-agentive concept of art, but in virtue of that artifact being *collectively represented* as an artwork. In other words, when function and procedure “become at odds”, an *artwork status-function* emerges and is imposed on objects, independently of these objects satisfying the causal-agentive concept. According to this description, we get two different and equally legitimate art-concepts: a causal-agentive concept of art and a status-functional concept of art. These concepts have different

extensions. For instance, the former includes objects produced outside the boundaries of the artworld institutions and the latter excludes them. Though they are equally legitimate and useful, the causal-agentive concept is more basic (explanatorily prior) than the status-functional concept. The latter depends for its existence on the former; it is a development or overgrowth. And the reason this is so is the simple fact that art *admits* of a causal-agentive concept. Artifacts which are primarily status-functional (e.g. languages, games, and other institutions) do not admit of causal-agentive variants, though the reverse is not true: some status-functions emerge out of causal-agentive functions and acquire a life of their own, so to speak (function and procedure “become at odds”).

To arrive at an understanding of how this is so, I would now like to focus on a particular kind of status-functional artifact, whose contrast with the category of artworks will (hopefully) prove enlightening: *games*. Games are status-functional artifacts because games are individuated by the rules that constitute them – whether *x* counts as an example of game *G* depends on what rules one must follow in order to play *x*. Whether a particular playing counts as a playing of chess depends on whether the rules of chess are being followed. The use of a physical chessboard and chess pieces is not necessary, as is made clear by the practice of blindfold chess; it is sufficient that players can relay chess moves between them. There are no chess moves independently of the rules of chess (the rules of chess are constitutive). Chess moves are powers that pertain to pieces in virtue of the particular pieces they are (pawn, knight, bishop, rook, etc.) and being a particular piece is not a matter of having a certain shape or constitution (e.g. being a miniature rook-wise piece of wood) – one can play chess using a computer, in which case the pieces are only pixels on a screen, or, more accurately, the pixel images are *representations* of pieces. But wooden, plastic or metal pieces on tangible chessboards are as much representations as are the images on the computer screen, and what is represented is a kind of power (to move, according with the rules). A chess piece is itself a status-function: *X* (a certain-shaped wooden, plastic, metal, pixel or purely imagined piece) counts as *Y* (pawn, knight, bishop, rook, etc.) in context *C* (following the rules of chess) and pieces are distinguished from each other by their powers to move. Chess pieces are thus defined by constitutive rules. Rules are status-functional artifacts: rules are individuated by how they are applied, i.e., by what it takes for a particular action to count as a following of a certain rule (action *A* counts as following rule *R* in context *C*). Therefore, to the extent that all games are constituted (and not merely regulated) by rule following, games are status-functional artifacts. What makes it

the case that the game of chess *exists* is that the constitutive rules of chess are collectively *shared* or *represented* as rules in use.⁹⁵

Famously, Wittgenstein is thought to have shown that no correct definition of *game* could be given, since (supposedly) no feature or group of features common to all games, in virtue of which they are games, could be found:

“Consider for example the proceedings that we call “games”. I mean board-games, card-games, ball-games, Olympic Games, and so on. What is common to them all? Don’t say: “There *must* be something common, or they would not be called ‘games’ “ – but look and see whether there is anything common to all. – For if you look at them you will not see something that is common to *all*, but similarities, relationships, and a whole series of them at that.”⁹⁶

In the late fifties, philosophers like Morris Weitz set out to apply Wittgenstein’s point about games to the philosophy of art, which made anti-essentialism about art more or less the default theory for the next decade or so. Then philosophers like Danto and Dickie, inspired by Maurice Mandelbaum’s critique of Weitz in his 1965 paper, “Family Resemblances and Generalization Concerning the Arts” and by art experiments with “indiscernible duplicates” (namely, Duchamp’s ready-mades and Warhol’s *Brillo Boxes*) went on to argue that the defining properties of art are perceptually indiscernible relational properties, such as an object’s relation to *art theories* (Danto) or an object’s relation to *artworld institutions*. Theories of this kind became known as *proceduralist* since Stephen Davies’ book, *Definitions of Art*, where he offers a classification of extant theories of art into three broad families: a) functionalist theories, b) proceduralist theories and c) historical theories.

Now, one point seems clear: if the unifying property of games were the fact that we call them “games” – the fact that we apply this noise/squiggle, “games”, to them – that would truly

⁹⁵ One complication arises if we think of chess (or any other game) as being invented at one time by a single individual. Should we say that the game does not exist until its constitutive rules are shared by several individuals? This seems weird. Surely the game exists from the moment it is invented, whether this is achieved by a single individual at one time or by several individuals over a long period of time. One way to counter this objection is to say that rules are intrinsically *shareable*, similarly to the way that chess moves are intrinsically *of* chess pieces. There is an analogy to be made here with words: though I can conceive of many phonetic/graphemic sequences that are not actually used as words in English as possible word-types in English, no sequence of phonemes/graphemes actually *counts* as a word in English unless it is actually used as such (shared) by speakers of English. One can also point to the difference between merely possible variants of chess that have never actually been played, and those variants of chess that have been (e.g. Tamerlane Chess or Shatranj): there is a sense in which the latter exist but not the former. The variants of the game that actually exist are those that figure in historical or causally connected chains of uses or followings of the rules that individuate those variants.

⁹⁶ Wittgenstein, 1986, §66 (trans. G. E. M. Anscombe).

correspond to what Stephen Davies has termed “a merely nominal essence”, but an essence nonetheless – it is true in all possible circumstances that the class of things we call “games” is only *that* class to the extent that we call them “games”. But this is a rather uninteresting property. It is one thing to define the class of things we call “games”, and quite another to ask whether among the things we call “games” there is an interestingly unified category of things (despite other extended uses of the concept). If a given class of things is merely unified by the fact that we apply some word to them, why should we still consider it a unified class at all? To take another example, it is (conventionally) arbitrary that we apply the word “chairs” to a certain class of objects. But there is a motivation for that practice: objects we call “chairs” share a core of properties in which we are interested (chairs are useful); and though some chairs cannot be used as seating devices (e.g. an art chair specifically designed to be unusable, displayed in a gallery), there must be some conceptual connection that makes it intelligible to call them “chairs” (unless, of course, the purpose is to make a witty gesture or to baffle the audience, in which case the conceptual *disconnectedness* is part of the point). Chairs can have radically different designs, but they will be part of a unified kind as long as they serve their primary function of seating devices. Dysfunctional chairs, those things that are unable to serve their purpose, though not as a result of an intentional action of their maker, will be chairs nonetheless, since serving the primary function of a seating device is what they *should* do. Functional properties are *normative*, i.e., the fact that an item X has the function F, and F consists in making it the case that X (or some other item by interaction with X) has properties A, B and C, does not imply that having (or generating) A, B and C is essential to X. It is not a necessary condition of being an artifact to actually have the capacity of discharging its functions. Otherwise there could be no dysfunctional chairs, since a dysfunctional chair would no longer be a chair. What this means is that an object *counts as* a chair to the extent that it has at least a *disposition* to perform that function *and* it was intended by its maker to do so. The condition of having at least some dispositional property that accords with the function is meant to exclude cases where the maker’s action in producing the object is so wide of the mark that the resulting object is not even recognizable as an attempt to make a chair (or take the example of the amateur who sits at the potter’s wheel with the intention of making a vase, but succeeds only in making a mess). Intentions are not enough; otherwise I’d be able to perform Bach’s fugued parts in the *Sinfonia* from the *Partita in C minor* merely by intending to perform them, while making nothing but some obnoxious percussive sounds at the piano. In the case of chairs, vases and similar artifacts, the intentions must be connected to an appropriate *insight* of the agent, about the kinds of physical properties that will sustain the relevant functional properties. Intentions are necessary, since any given artifact will be able to perform functions other than its primary function, but that gives us no base to qualify, say, a chair as a

climbing device, though it can be used to perform what would be the primary function of, say, a small ladder. Likewise, we can use a stone to hammer a nail, without the stone thereby qualifying as a hammer. The point is that the functional properties of artifacts (which are their defining properties) are *historical* properties – the kind of property that the “relational” theorists of art thought would defeat the standard objections to the definition of art: non-perceptual relational properties.

This already makes apparent one mistaken assumption in proceduralist theories: their original proponents assumed that a “relational” definition of art will disqualify traditional functionalist theories in favor of the kind of theory held by Danto and Dickie. But the only kind of traditional theories of art excluded by “indiscernible duplicates” are non-normative essentialist theories, such as a theory according to which artworks have aesthetic properties *essentially*.⁹⁷ Such theories in fact could not account for the existence of *bad art* or dysfunctional artworks. But it is even doubtful that any proponent of functionalism has ever entertained such a view, though it is certainly conceivable.

The phrasal verb “counts as” will be of crucial importance in our understanding of the relation between theories of art and also between different kinds of artifacts. I’ve already specified two wholly distinct kinds of functional properties, a distinction based on the way an object is *able* to serve those functions: some functions are dependent on the intentions and insights of cognitive agents and do not require collective acceptance of a *status* in order to be of the functional kinds they are of – such properties give us a kind of artifactuality which I term “causal-agentive artifactuality” (a notion inspired in John Searle’s taxonomy of functions). They are so called because the objects have some (dispositional) power to perform a function, *and* that function was intended by a cognitive agent, based on an insight that some physical structure would be appropriate to sustain the relevant function. Some objects can be multifunctional by design (like the Swiss army-knife) and not by appropriation (the case in which an artifact with a certain primary function is *conscripted* to perform some other function – e.g. when a chair is used to climb). A different example of an artifact being conscripted to perform a function other than its primary function is when *words* are employed in figurative speech: their literal meanings are “appropriated” to perform a communicative function that goes beyond their literal meaning.

Words form a different kind of artifact from paperweights, hammers, and chairs: they are *status-functional* artifacts, since in their case the intentions and insights of a cognitive agent are

⁹⁷ In fact, this basic mistake can already be seen if we take in consideration the fact that the “family resemblance” account of the art concept, while being anti-definitional, is an account *of the nature of art* (of what in the nature of art makes the concept insusceptible of a classic definition) *and* it is an account in terms of a relational property that has nothing to do with institutions – resemblance.

insufficient to make it the case that the norms for artifactuality are successfully brought to bear on them. To think otherwise would be to embrace the private language hypothesis. But this won't work. It is in virtue of this difference in kinds of artifactuality that "humpty-dumptying" defeats the purpose of word use. If individual intentions were sufficient to give meanings to words, the use of words would be pointless. Words are artifacts that have their functions in virtue of *convergences in behavior*; they are *conventional* artifacts, though here conventions come in two distinct brands: some words are *stipulated* (e.g. "quark" was introduced in 1963 by Gell-Mann to refer to a kind of physical entity, though it had already occurred in James Joyce's *Finnegan's Wake* – this was both a case of appropriation and stipulation); but the vast majority of words was not brought into being by stipulation. Words are primarily *unintended emergent artifacts* (the phrase is Zangwill's), but their artifactuality is status-functional; nothing is a word, nor the particular word it is, apart from shared representations and a convergence in use. If I use the word "music" with its current meaning – as a kind of sonic phenomenon – and you use the same word with its medieval meaning – something that had sonic "manifestations" but was primarily a non-aural phenomenon – we are both using (overlapping) *concepts* of music, equally articulated by the word "music", but we are not using the same concept. I'll be referring to an art form while you'll be referring to a mystical conception of "cosmic order" that is faintly mirrored by *some* human art music (curiously enough, the Pythagorean concept of *music* is, in face of our current use of the term, a *metaphorical* use, while the current use is literal, even in terms of the ancient and medieval concept). But *what* concept each one of us is using does not depend merely on our individual intentions, but on what conventions are at work in the way we use words.

Going back to games, the point about extended and figurative uses should immediately alert us to a methodological deficiency in the "family resemblance" approach: we cannot pack together those objects to which the term "game" is applied in an extended or figurative sense and the objects to which it is applied in its literal, primary sense, and then be surprised that they don't form a unified kind. When one says of his dog's familiar barking that "it is music to my ears", one is using the term "music" in a non-literal sense. So that pointing out that use cannot give one a reason to claim that "music" does not refer to a unified kind. And to distinguish between extended, figurative and literal uses, one must do more than "look and see". A good amount of reasoning must be involved. Furthermore, the non-literal uses *depend* on the literal use. When I utter the metaphorical sentence "No man is an island", the identity of the metaphor depends on the literal meaning of the words "man" and "island". If they had a different meaning, I would not be making the same metaphor; I might not be making a metaphor at all. What *I* mean by that sentence is something like "No human being is self-sufficient" or any other appropriate paraphrase. But "island" is not being

used with the meaning “self-sufficient entity”; it is being used with its literal, primary sense. Or take the metaphor “metaphor is the key to knowledge”, uttered to make a point something like “The way we represent the world would be severely handicapped if we lacked the ability to explore conceptual mappings between disparate domains of objects” or “figurative language allows us to represent things and states of affairs we would not be otherwise able to represent”. It would be ludicrous to present such usage of the word “key” as evidence that the class of things that are keys is a “family resemblance” class, or that it is hopelessly gerrymandered. No metaphor is literally a key, but the metaphor is intelligible because *what we mean* by uttering the figurative sentence is conceptually dependent on our previous grasp of the concept *key* in its primary, literal use. We play with the concept of “having the power of opening locks and allowing access to a space and what it contains”. Keys have this property literally, as part of what defines them as the functional entities (artifacts) they are. A more mundane figurative use of the same concept is when we say that “X is of key importance to understand Y”.

So, our methodology to understand the nature of games must go beyond the “look and see” precept. We must ask whether among the things we refer to as “games” there is a unified kind of thing that makes extended uses of the concept intelligible. One example of such extended use is when we say things like “philosophical discussions are nothing but a game”. Just like in the key example, where we play with the property that keys have of opening locks and allowing access to a certain space, in this latter example we play with a relevant property of games, one which we may find difficult to articulate explicitly but one to which we intuitively “triangulate” in normal language use, and makes sentences like “this is a mere game” (said of things that we do not consider games in the same sense that chess is a game) intelligible.

Jim Stone, in a 1994 paper titled “Games and Family Resemblances”, has explored these issues, and proposed a definition of “game” that attempts to illuminate such extended uses of the word. He rejects the methodology of looking while avoiding thinking – “Wittgenstein’s injunction, ‘Don’t think, but look!’ is an invitation to miss the forest for the trees.”⁹⁸ In his view, “A game is a rule-defined activity involving a state which counts as performing the activity successfully because it is so defined by an arbitrary rule, an activity typically performed for the recreation of participants or spectators, or to sharpen skills.”⁹⁹ So, take the case of chess: what are players trying to do when they play chess? They are trying to reach the “state of success” (which can either be achieved by checkmating the opponent or if he resigns). And *why* are they trying to do this? Because the rules

⁹⁸ Stone (1994: pp. 436-437).

⁹⁹ *Ibid.* p. 437.

specify such a state as the state of success. Players do not strive to reach the state of success because that is a desirable outcome independently of the constitutive rules of chess. The rules *make it the case* that such and such a state is the state of success. This is a case of what Searle called a “status-function declaration”. (Stone’s paper preceded Searle’s first book on social ontology and it is curious how many things that Searle says are anticipated there; it is immensely interesting to read Stone’s article *after* having read Searle on this subject.) The rules of chess differ from the convention of driving on the right (or the left) in the sense that the rules of chess do not merely *regulate* some previously existing activity that is independent of the existence of traffic conventions. People drive on the road because they are pursuing a goal which is desirable independently of any arbitrary rule’s defining it as a state of success: they want to reach their destination. But chess players are not engaged in the same kind of activity, though both activities are actually governed by rules – the rules of chess in one case, the traffic conventions on the other. Of course, the reasons for playing chess are many. One good reason is that it is *fun* to do so; another is that it is challenging, and exercising one’s intellectual skills is also a pleasurable activity. There is a specific pleasure that we get from playing chess that we cannot replace by other pleasures, e.g., that of listening to a piece of music, though listening to (and playing) music is also intellectually challenging and in part a pleasurable activity for that reason. No arbitrary rule determines the pleasure of listening to a symphony with understanding as a state of success. We listen to music because we find it pleasurable. We also play chess because we find it pleasurable (though one could also play because one has made a bet, or we can imagine fanciful counterfactual situations in which one is forced to play chess for some reason), but *what we are doing* while playing chess is to pursue a state of success which is such not because of the pleasure we get from it, but because an arbitrary rule stipulates it. It is the overcoming of the difficulty in reaching the state of success which is pleasurable. Checkmating the opponent doesn’t count as success because we strive for it; rather we strive to checkmate the opponent because it is success *according to the rules*. One can imagine cases where a game of chess is interrupted and no one reaches the state of success; or cases where the game ends in stalemate; but this does not impugn the fact that while the players are engaged in playing, what they are doing is trying to achieve the state of success. This is what *defines* what they are doing *as* an instance of playing chess; since nothing *counts* as playing chess independently of the arbitrarily stipulated state of success. All the other rules of chess would be pointless if they were not the establishment of boundaries in terms of which one tries to reach the state of success. Things

could have gone differently. Chess could have evolved to have different rules; something which is clearly attested in the existence of *variants* of the game.¹⁰⁰ The state of success could be differently defined. But none of these possibilities defeats the point that chess is defined in terms of its constitutive rules and that nothing could intelligibly count as an instance of playing *that game* if the players were not pursuing the state of success. Someone who shuffles chess pieces on a board without regard for the state of success is *not* playing chess, though what he does may be an instance of *playing with a chessboard and pieces*.

As with the case of “indiscernible duplicates” in the philosophy of art, we can imagine cases where an activity physically indiscernible from playing a game is *not* an instance of playing a game. Consider the case where someone throws a ball against a wall and catches it. I can do this for the mere fun of catching the ball, without any concern with whether I catch it a number of times in a row, with a particular way in which I catch the ball, etc. But what I do will count as an instance of playing a game if while doing it I entertain something like the following thought: “Catching the ball *n* times in a row without dropping it counts as success” or “Catching the ball *n* times in a row *in a certain manner* counts as success”. Two people, each engaged in an activity physically indiscernible from what the other is engaged in, can differ in that one counts as playing a game while the other doesn’t. Frame this as a Twin-Earth scenario, where someone on Earth throws and catches the ball exactly like his twin-earthling is doing on Twin-Earth, except that one is entertaining the kind of thoughts I mentioned, while the other is just throwing and catching the ball for the pleasure of doing so. One of them is trying to reach an arbitrarily defined state of success, and therefore is playing a game, while the other is not. It is not the case that there is some *vagueness* in the situation, to the effect that it is not clear whether one of them is playing a game and the other isn’t. There is no vagueness, even if an observer could not distinguish both cases merely by “looking and seeing”; just as there is no vagueness in the case where a set of wood *Brillo Boxes* is a case of art, while a set of carton *Brillo Boxes* is not art.

Another potential case of indiscernibility is that between a boxing match and a street brawl. Here we run up against one of Wittgenstein’s original examples of the gerrymanderedness of the class of games: “board games, card-games, ball-games, Olympic Games, and so on.” We count boxing among the Olympic Games, though we are ambivalent toward calling it, and other sports, “games”. But here a theory of the nature of games, though defeasible (as such theories are and should be),

¹⁰⁰ For a list of chess variants, see:

<http://www.chessvariants.com/index/mainquery.php?type=Any&orderby=LinkText&displayauthor=1&displayinventor=1&showimages=1&showtextnotes=1&listrecognizedonly=1&primarylinksonly=1&usetthisheading=Recognized+Chess+Variants>

turns out to be more enlightening than the “family resemblance” approach, which does nothing to explain our ambivalence. If games are connected by no more than “overlapping resemblances” why should we have a more ambivalent attitude towards some cases but not others? The “family resemblance” approach gives us the same mechanism – resemblance – to explain both our lack of ambivalence in some cases (we classify this as a game because it resembles other things called “games” in certain respects) and our ambivalence in other cases (we hesitate to call this a game because it resembles other things called “games” in certain respects but not all). This is explanatorily empty. Now, a definition like Stone’s illuminates the ambivalence in our attitudes: to think of boxing as a game, we must think that knocking out an opponent counts as winning because a rule says so (we can easily imagine a counterfactual situation where knocking out the opponent counts as losing; in fact, in a “game” situation, there are ways of knocking out one’s opponent that will not count as winning, the same way not all instances of getting a ball to a goal in soccer will count as a “score”; but the same idea makes no sense in the case of a street brawl: if you knockout your opponent in a street brawl, you win; the end). In a street brawl, which is *not* a rule governed activity, knocking out your opponent counts as winning because *it is* winning. So, the root of the ambivalence is that the state that the rule stipulates as the state of success *already is* the state of success in any instance of fighting which is not governed by rules. This accounts for our ambivalent attitude towards the example of “Olympic Games”.

This also illuminates certain extended uses of “game” as in sentences like “philosophy is just a game of words.” What is being said here is that people engaging in the said activity (philosophical discussion) behave as if they were pursuing an arbitrarily specified “state of success”, which counts as success *not because it is independently desirable or valuable* (e.g. because it would generate knowledge or insight, in which case it would count as success in virtue of *being* success, independently of what are the rules for “counting” here), but because the rules say so. The operative distinction behind “philosophy is just a game” is that between what is strived for in an activity in virtue of itself constituting the purpose of the activity (e.g. acquiring true beliefs), and what is strived for because that is what some rule stipulates (e.g. cowering the audience into an awed silence; appearing erudite, etc.). To arrive at the truth counts as success because truth is something we value or strive for, not because an arbitrary rule defines it as success. The point of the jeer is that people engaging in the said activity are (allegedly) not, say, pursuing truth (something that would independently count as success) but rather behave as if they were in fact playing a game, that is, pursuing an outcome that counts as success because that is what the rules say it is. When someone has “won” the argument, what he did may involve pleasure and the exercise of intellectual abilities, but what he has achieved, so our jeerer infers, is not independently valuable, like a driver’s goal of

reaching his destination, which can be done in a situation where there are traffic rules, but remains the same goal in a situation where there are no traffic rules.

Consider one more example given by Stone. Waltzing isn't a game according to the definition. "To waltz successfully, it is, of course, necessary to waltz. But that isn't sufficient. Successful waltzing is waltzing with grace, ease, and beauty."¹⁰¹ This is actually misleading. To *dance* one must move in a certain way that is distinct from the way we move in everyday affairs. Moving with grace, ease and beauty is one cluster of properties that can be sufficient, but dancing can occur with movements that exemplify different aesthetic properties. Waltzing introduces a complication and clarifying it will prove useful for the way we are to conceive of art. Waltzing incorporates a NKA – dancing – an activity whose "state of success" is given by the definition of the activity independently of any regulative rules that we incorporate in dancing (by incorporating regulative rules, we add "layers of convention" to the basic NKA, which results in a "culturally modulated NKA"). In this sense, waltzing is not a game, because it crucially depends on a NKA. But what *counts as* waltzing is not *simply* given by the nature of dancing; something else is required. There is a parallel here with *impressionistic depiction* – depiction is a NKA (assuming Goodman is wrong and "perception theories" of depiction are right), but what counts as *impressionistic* is not simply given by the nature of depiction. Succeeding at depiction is not sufficient for succeeding at impressionistic depiction. To produce an impressionistic depiction one must bring the conventions of impressionism to bear on what one does. This bringing to bear can be an originating use – as was the case with the first impressionist paintings – or it can be a deferent use – as was the case with later impressionist paintings. But this is relevantly different from the case of games, because what counts as success in producing an impressionist depiction does not depend on a *constitutive convention* (like the rules of chess are constitutive conventions). Success at waltzing, at producing impressionistic depictions, like success in making a symphony, a movement in sonata form, a fugue, a performance in the tradition of Persian *radif*, etc., involve specific conventions (such the conventions about sonata-form movements) that regulate music making (e.g. what counts as making a symphony in the early nineteenth century differs from what counts as making a *sinfonia* in the early eighteenth century).

Stone's definition is preceded by the relatively "unknown classic" in the literature on games, namely, Bernard Suits' *The Grasshopper: Games Life and Utopia*, originally published in 1978.¹⁰² Here, Suits defends a definition of the action of game playing (rather than of "games" as a class of

¹⁰¹ Stone (1994: p. 442).

¹⁰² Suits (2014): Ch. 3; statement of the definition of games in p. 43.

objects separate from a specific kind of *action* of playing) that hinges on three proposed necessary and sufficient conditions: 1) a *prelusive goal* (another name for “state of success”) or specific state-of-affairs which one can describe independently of the game; 2) *constitutive rules* which prohibit the use of the most efficient means in favor of less efficient means (according to Suits, it is characteristic of games that they prescribe the surpassing of “unnecessary obstacles” because human agents value such pursuits as ends in themselves); and 3) a *lusory attitude* or the acceptance of the constitutive rules in virtue of the fact that they make the activity possible. This is the same as Stone’s notion that one strives to reach the “state of success” not because it is success independently of the rules, but because a constitutive rule arbitrarily *stipulates* that such state counts as success. The means-limiting aspect of constitutive rules is thus inextricably bound to the arbitrariness of the “state of success” in game activities.

The definition of *game* proposed here also excludes some things commonly classified as games (which is only to be expected of any interesting definition – a definition which doesn’t exclude anything is *not* a definition). An example given by Stone is that of a brand of video “games” titled *Sims*, where the player observes the “lives” of characters he creates in a simulation, but there is no “success state” arbitrarily defined. The “game” actually never ends, unless the player shuts the program off. This does not count as a counterexample to the theory. It looks like a counterexample because the program is made to run in a digital computer and is sold as a “game” and stored with other computer programs that *are* games. It also bears several similarities to computer games, but in fact is nothing but a “virtual” version of playing with dolls, something that we more intuitively disqualify as a game. Computer games count as games when playing them consists in trying to reach an arbitrarily defined state of success.

The “Don’t think but look!” methodology turns any concept into an arbitrary concept, since it automatically counts “confused cases” (such as virtual doll houses sold as “games”) as counterexamples to any definition. Since we can’t afford to be led astray by thinking too much about what we see, no attempt at understanding *what* is confused in a “confused case” will be considered a legitimate move by the “family resemblance” theorist, which means we won’t be able to illuminate the differences between uses of a concept in a primary, literal use, and uses that are extended or figurative; or at least we won’t be able to do so while remaining faithful to the “family resemblance” doctrine. This view construes practically all concepts as arbitrary concepts. But arbitrary concepts are inoperative and uninteresting. We have no use for them. We want *explanation* and clarity; but the “family resemblance” doctrine does not illuminate, it merely adds to the confusion. This was already substantially addressed in the literature: resemblances are cheap! Any two objects will resemble each other in *some* respect, so that if the “family resemblance” view of

concepts is to be of any use to us, it must provide an account of *what* resemblances are relevant and what resemblances are not. But to do this is already to eschew the “Don’t think but look!” methodology. Even if the “family resemblance” doctrine is right in the way it describes how we acquire and use certain concepts – by way of salient similarities (whatever it is that makes them salient) rather than by the grasping of some “essence” mysteriously “contained” in the word or concept – this does not defeat any definitional approach to art, games or any other social and artifactual kind; since the truth of that *epistemological* point is compatible with the *metaphysical* one I’ve made about there being an interestingly unified *kind of thing* referred to by a social or artifactual kind term, that makes extended and figurative uses of the same term intelligible. Of course, we can arbitrarily *decide* to apply the word “games” to whatever we like – Weitz’s point about art being an open concept amounts to no more (or little more) than this. But this is *irrelevant*. Regardless of how the conventions for the use of “game” vary, there is an interestingly unified kind of thing, which has the properties specified by Stone’s treatment of games. What we must investigate is whether the existence of this kind of thing, and our ability to grasp it as a unified kind of thing, even if only indirectly or confusedly – as the cases like “Philosophical discussions are nothing but a game of words” suggest – is what makes the concept and the distinction between its primary and extended or figurative uses *intelligible*. The Pythagoreans and those medieval thinkers inspired by Platonism and mysticism had a concept of *music* wholly different from our current concept. But what is relevant is not the *possibility* of such wide divergence in conventions for word usage. What is relevant is whether, among the things we call “music” there is an interestingly unified kind of *thing*, in particular an interestingly unified kind of *experience of sounds*, that actually throws light on what superficially looks like a chaotic mesh of uses of a certain noise – “music”. The same point, of course, applies to *art* and to any social and artifactual kind.

This is not to say that theoretical approaches such as that of Stone cannot be *wrong*; it is not to say that they can’t be substantially revised, clarified, refined, etc. But I’ll take an illuminating mistake (a mistake that makes you think) over a thought-stopper (though apparently “revelatory”) injunction any day of the week. In this way, one can look at the history of contested theories and definitions of art not as wreckages that attest to the barrenness of the pursuit, but as tentative progressions in triangulating to art’s defining properties, not because achieving a definition free of counterexamples is something like an arbitrary success state, that is good insofar as it makes us exercise our intellectual skills in the process, but because each progress, tentative and defeasible as it may be, generates *explanation*; it throws light on the subject and dispels confusion. This is why I don’t see the abundance of theoretical approaches to the ontology of art as “an embarrassment of riches” but as a welcome state of affairs. Any theories about the nature of artworks may all be wrong, but going

through the details of such “embarrassment of riches” pays off, in terms of how one’s thought about art is *clarified* in doing so, even if one ends up rejecting a given theory. We may not have reached the knowledge we expected, but that does not mean that we don’t know *more* partly *because* of that “embarrassment of riches” fueling our curiosity. In a reply to Gary Iseminger on the role of audiences in a theory of art, Nick Zangwill takes “the variety of options for an aesthetic view of art to be a testament to vitality of that tradition.”¹⁰³ This is exactly my attitude towards the variety of options in art ontology.

Enough about the definition of *game*. The point is not so much to prove Stone’s definition, or anyone else’s, true, but in part to break the varnish of unmovable dogma that has been poured onto the anti-essentialist stance. Also, games are an important example of status-functional artifacts. Analogously to words, games depend on shared representations (the rules). The fact that moving a certain piece in such and such a way *counts as* a legal move in chess depends on a convergence in representation. If game rules depended solely on individual intentions, there would be no games, just as words become pointless in a world where everyone engages in systematic “humpty-dumptying”. More than that, humpty-dumptying, like the figurative use of words, can only be conceived in a world where words have publicly established primary, literal meanings. The same is true of games: people can disagree about the rules because there *are* publicly established rules for the game in question. The fact that there are publicly established rules is also what establishes the *existence* of the game. Games would not be possible without the explicit or tacit use of the phrasal verb “counts as”, which in turn makes games, like all institutional entities, language-dependent, since only beings endowed with a language (in the robust sense of something that has the properties of discreteness (repeatable elements, like phonemes, words and sentences), compositionality and a generative syntax) can represent something *counting as* something else. A non-linguistic animal can see a ball hitting the goal, but it cannot see a *score*, since to see that one must be able to *count* balls hitting goals *as* scores. A social group without language can have alpha-males or alpha-females, but this is not an institution. We have a case of institutional reality when people *represent* someone *as* being the leader (or the priest, or whatever), and that person has certain (positive and negative) powers not in virtue of his continued ability to beat up everyone else, but in virtue of the fact that the others represent him as endowed with a certain status. This is the general form of a status-function: “X counts as Y in context C”. It is the *form* of the shared representations on which status-functional artifactuality depends, even if those representations are never actually *expressed* in

¹⁰³ Zangwill (2009: p. 39).

explicit sentences of that form. One doesn't have to "think out loud" that a ball hitting the goal in such and such a context counts as a score in order for it to be the case that our shared representations determine when a certain physical event counts as a score in a game.

The previous point about games mirrors the structure of institutional facts more generally. Take the rules for electing someone president of an organization. We settle on a (conventionally arbitrary) rule according to which the person who gets more votes, given the observance of certain contextual requirements, counts as president. We could equally as well settle for a rule according to which the candidate that manages to roll down a hill and reaching the foot of the hill with the largest amount of grass in his mouth counts as having won the presidency. That would be awkward or bizarre given our cultural context, and perhaps impractical, but surely not inconsistent. The point is that the mark of the institutional is that the state of success is strived for not because it is success independently of the rules, but because the rules fix what counts as the state of success. In Stone's words: "succeeding [in games] is wholly internal to the game, which is what defines success." The same applies to institutions generally: when a social practice constitutes an institution, what counts as success in the practice is wholly internal to the "game"; the game defines success, not the other way around – which is the case when a practice falls short of being *institutional* (like in the street brawl and the waltzing examples). This is important because it gives a neat distinction between *social practices* and *institutions*, a distinction which many philosophers of art fail to make, which in turn can have important consequences for how, e.g., institutional theorists conceive of art in institutional terms. One example I have in mind is Gary Iseminger's (2004) aesthetic-institutional theory, which, as I do in this work, draws on Searle's views about social ontology, but arrives at a conclusion I reject, namely, that art practice is an institutional practice and artwork status is a matter of status-functional artifactuality. (Another theorist who thinks of artwork status in terms of status-functional artifactuality is Peter Lamarque¹⁰⁴. A similar view is held by Catharine Abell¹⁰⁵ who also draws on Searle's theory of institutional reality.)

Iseminger's theory is an aesthetic theory to the extent that it defines the *role* or function of the "artworld" (which is a set of social practices thought of as an institution) as the "promotion of

¹⁰⁴ Lamarque (2010: p. 69) "The continued existence of any work depends on the continued possibility of the work's being *responded to* in appropriate ways. A work is sustained in existence partly in virtue of the attitudes, beliefs, and desires of those who recognize its role as a work and as the work it is. (...) It is not an 'ideal' entity, existing only in some mind (that was one of our desiderata), but it is an *intentional* entity, depending essentially on facts about how it is *taken to be* by qualified observers." – At the end of the first sentence, Lamarque opens a footnote quoting a passage of Searle (1995) where he explains how the status of institutional facts is constituted by collective acceptance. Lamarque suggests that artworks "share common features with other institutional objects or facts."

¹⁰⁵ Abell (2011: p. 14) "On the approach I am advocating, the notion of an artwork is essentially institutional, because facts concerning which things are artworks are institutional facts."

aesthetic communication”. The theory doesn’t define *artwork* directly, but rather indirectly – artworks are those things that acquire artwork status within the relevant practices. What Iseminger does is to define the *artworld* (the set of art-related social practices) as a kind of (status-functional) artifact, one whose proper function is to promote what he calls “aesthetic communication” (it is not important for our purposes here to delve into the specifics of this concept). According to the theory, the artworld has this function contingently (things could have come to pass in such a way that the artworld had a different proper function). Being an artist, within the relevant social practice, is thought of in terms of a status-function, someone *counts as* an artist in a similar way that someone counts as being president, though there is a difference in that status conferral within the artworld lacks the formality of other institutional practices. This is a relevant point, in that George Dickie, the main proponent of an institutional theory of art, was lead to revise the clause about status conferral in his first definition by criticisms to the effect that art practices are informal, which seemed to exclude such a stately and pompous thing as status conferral from the world of art. This historical curiosity already shows how a failure to get one’s social ontology minimally right may interfere with the prospects of arriving at a sound theory of art kinds, institutional or otherwise – since non-institutional theorists must also have their account of *why* art is *not* primarily institutional. Stephen Davies has (rightly on my view) criticized Dickie for abandoning the clause on status conferral, on the grounds that status conferral is essential for the theory to be properly called “institutional”, rather than just a “necessity of social context” theory (or in Dickie’s own terms, a theory of the “necessary and sufficient cultural matrix” of artworks).¹⁰⁶ Davies got his social ontology minimally right (whether he was then right to endorse a version of the institutional theory is a different matter).

We can distinguish institutional theories (or would-be institutional theories) from (non-hybrid) aesthetic theories by the use they make of the phrasal verb “counts as”. For (non-hybrid) aesthetic theories (which I subsume under “causal-agentive artifactuality theories”) an object counts as an artwork because *it is* an artwork, independently of what any (conventionally arbitrary) rule specifies as success in status conferral. For such theories, success in making art is not “wholly internal” to the “artworld game”. Success in art making is like success in waltzing or like success in a street brawl. Whether we specify such success with the concept ART or any other concept is not of utmost relevance. We need *some* concept to refer to that (language-independent) state of success. Suppose that aesthetic functionalism about art gets it right: that state of success will be specified by the

¹⁰⁶ Dickie (1997: p. 27).

concept of producing an artifact with certain causal-agentive functions, namely, those of having certain aesthetic properties in virtue of having certain non-aesthetic ones. We need a word to refer to this state and “art” is just fine. We could have had any other. Nothing of special importance hangs on this.

For institutionalist theories, an object counts as an artwork because it is embedded in a system of shared representations that have the basic form “X counts as Y in C” (e.g. the snow shovel counts as an artwork in the twentieth century artworld, etc.). What makes institutional theories *institutional* is the fact that success in art making is thought here in terms of a “game” (the artworld game) to which succeeding is “wholly internal”, that is, not independently specified by the *nature* of the activity before the “rules” were brought into play. But not all proceduralist theories are institutional (though all institutional theories are proceduralist); and now we have the tools to see clearly why this is so: not all (social) procedures are institutional. Danto’s earlier theory of artworks as things that are so in virtue of a relation to *art theories*, and Dickie’s later theory of artworks as artifacts of a kind to be presented to an audience, are of this kind – procedural but not institutional in the strong sense. Dickie’s later theory has the further disadvantage that it can be made compatible with an aesthetic theory, as long as the aesthetic theorist is willing to grant a “necessity of social context” clause; basically, it amounts to the recognition that art is a cultural activity, which no aesthetic theorist has trouble accepting, since many cultural activities are but “culturally modulated natural-kind activities”, that is, the fact that they are culturally inflected is compatible with “success” in those activities *not* being a matter of status-function rules (in which case what Dickie calls a CKTA would be called for, and not a NKTA). Take once again the case of depiction. Goodman’s conventionalism is a “cultural-kind theory” of depiction. Here depiction is not thought of as a “culturally modulated natural-kind activity”, but a “game”, with a state of success internally defined by the rules – if you don’t know the rules, you can’t understand depiction. For such a theory, “successful” depiction is a matter of denotation, not perception (thus, status-functional, not causal-agentive). Theorists like Richard Wollheim, on the other hand, have a “natural-kind theory” of depiction. Here depiction is not conventional, though it can incorporate layers of convention (it can be and probably always is “culturally inflected”). Success in depiction, for “seeing-in” theories such as Wollheim’s, is not defined by a contingent rule; it is specified by the nature of the perceptual experience of seeing something (the depicted item), e.g., in a marked surface. (See e.g. Wollheim (1980: pp. 205-226)).

In a supplementary essay to *Art and Its Objects*, where Wollheim contrasts the notions of *seeing-as* (seeing X as F) and *seeing-in* (seeing F in X), we find the following remark:

“... seeing-in derives from a special perceptual capacity, which presupposes, but is something over and above, straightforward perception. This special perceptual capacity is something which some animals may share with us but almost certainly most don’t, and it allows us to have perceptual experiences of things that are not present to the senses: that is to say, both of things that are absent and also of things that are non-existent.”¹⁰⁷

Examples of such experiences of things “not present to the senses” and of “things that are non-existent” are found in the discerning of scenes of violent action, mysterious landscapes and other events or objects when one looks at “damp-stained walls or at stones or broken color”.¹⁰⁸ The experience is perceptual through and through, though it is also something “over and above straightforward perception”, namely, it differs from the latter in that it is characterized by “twofoldness”, i.e., in it, we are both perceptually aware of the features of medium, the marked surfaces, and the depicted object (that which we see *in* the marked surfaces).

The parallel with depiction must be emphasized. Although there is a practically universal consensus that not all representation is conventional – namely, depiction is not, at its core, conventional, though it can be (and probably always is) culturally inflected – dissent is possible. Goodman defended conventionalism about *all* representation, including figurative depiction. He also defended that the arts belong to “symbol systems” (so, as in Danto, art is here conceived as “a language of sorts”) and this accords with Wittgenstein’s conception of language as a system of “language games”. The theory of games just explained helps us to make better sense of this fundamental idea: Wittgenstein opposed the Augustinian view of language as a system of signs (sonic or graphic) with “derived intentionality”. On the Augustinian model, humans have a *natural* ability to think about things and to add language to this picture consists merely in the transfer of intentionality from thoughts to words. To speak a language is thereby conceived in the same manner as a natural-kind activity, such as walking, brawling and dancing (and seeing depicted objects in marked surfaces, if one holds a “perception theory” of depiction). The picture of language as a collection of language games, on the other hand, opposes this view of language as a NKA. Words only have meaning as determined by their use in a “language game”; this is what is implied in words being *status-functional artifacts*, not causal-agentive artifacts (as they are thought of in the Augustinian model). And the rules that determine how the game is played have no “derived intentionality” in virtue of which they can be applied; the meaning of a rule is determined by how it

¹⁰⁷ Wollheim, “Seeing-as, seeing-in and pictorial representation”, in 1980, p. 217.

¹⁰⁸ *Op. cit.* p. 218.

is applied, not the reverse.¹⁰⁹ If language is thought of as a game (as status-functional) then meaning is determined by use, as is to be expected if Wittgenstein is right. Thus the theory of games I'm drawing on throws more light on the concept of language games than the "family resemblance" view of concepts does. It also throws more light on the idea of institutional reality consisting in a system of horizontally interlocked and vertically iterative status-functions than Searle's own internalist view of language. Furthermore, it accords with a consequence of everything I have said about kinds of functional properties and theories of art: causal-agentive theories (such as aesthetic functionalism about art) see art at bottom as language-independent (though many conventional aspects that art integrates will be language dependent – e.g. humans without language can have art behaviors or music behaviors in this model, but they can't have symphonies, movements in sonata-form, and all that – these are only possible when *natural* art behaviors and music behaviors integrate regulative conventions, and it is these that are language-dependent). On the other hand, status-functional theories (such as institutionalism about art), see art at bottom as a language-dependent activity. This is what is to be expected if philosophers like Dickie are right about the practice of art requiring a conception of art. If art practice was as the institutionalist theory conceives it, *succeeding* in that practice would be determined by the "rules"; making artifacts that sustain aesthetic properties according to insights and intentions of their makers (causal-agentive artifactuality) would not be sufficient or necessary. Making something endowed with art-status, i.e. succeeding in the "artworld game", would not be supervenient on the nature of an activity that can exist without regulative conventions; it would supervene on the constitutive rules that make the "artworld game" possible (if art is to be conceived as having the same matrix as games). In sum, if philosophers like Goodman, Danto, Dickie and Iseminger are right, then art is more like language and games than like walking, brawling, dancing and depiction (as "perception theories" conceive of it). If philosophers like Beardsley and Zangwill are right, then art is more like the latter than the former.

We can now introduce a further distinction into Dickie's classification of theories of art into NKTA and CKTA: the latter group will include "social contextualization" theories, such as Danto's theory of 1964 and Dickie's later version of the "institutional" theory, of 1984; and properly so called *institutional* theories, such as Dickie's earlier version of his theory (with the status conferral clause); Davies' own view of the institutional theory (which has no explicit definition to accompany it); and Iseminger's hybrid theory. Institutional theories proper share a fundamental structure with

¹⁰⁹ Stone (1994: p. 440).

games, according to the theory of games I've explained above. The (proposed) definitional property of games gives us the matrix of institutional practices: all institutional practices are social, but not all social practices are institutional. To be institutional, a social practice must have some of the properties that are definitional of games: a) it must be defined by *constitutive conventions* (as these were explained in 1.2) – conventions that do not merely regulate a preexisting activity – and b) what counts as *succeeding* in such practices is not specified by the nature of a social activity that can exist without regulative conventions (e.g. walking, driving, brawling, dancing, etc.) but is rather specified by the rules that make the activity possible.

Iseminger's aesthetic-institutional theory is in fact an institutional theory (unlike Dickie's later theory), but I believe it ultimately fails, because though I grant that we have a perfectly sound concept of arthood as a status-function – there is such a thing as “artworld art” that excludes much that is not excluded by an aesthetic theory of art while including much that the latter would exclude – I don't think that art is *primarily* status-functional. Stephen Davies points in the right direction when he says that functionalism about art (what I call “causal-agentive theories”) and proceduralism (what I call “status-functional theories”) regularly beg the question against each other, in that theorists in both traditions think that art must be *either* functional (causal-agentive) *or* procedural (status-functional), where the disjunction is an exclusive disjunction. Having endorsed art concept pluralism, the thesis that there isn't a *single* consistent art concept (see previous section), I find it plausible that we actually have more than one perfectly sound concept of art; “art” is ambiguous between causal-agentive art and status-functional art. But the first concept has metaphysical primacy over the latter. Once more, Stephen Davies points in the right direction when he argues that we must acknowledge an important functional dimension to art if we are to make sense of how the arts and the “artworld game” have arisen. First art, in his view, must be functional. It is because we experience the world aesthetically and are driven by the impulse to explore aesthetic properties in the making of artifacts that an art practice has emerged. The first artistic artifacts are objects that were made to instantiate aesthetic properties (properties such as an appealing symmetry, balance, delicacy, elegance, and so on)¹¹⁰ even if they are also (or primarily) made to serve other functions. Here Davies agrees with the aesthetic functionalist. But as language

¹¹⁰ The case of Acheulean hand axes, stone implements of around 400kya: elongated, leaf or tear-shaped, symmetrical blades that show little or no signs of wear. About them, Gregory Currie (2009: p. 12) has remarked: “In size and shape it would not have been a useful butchery implement, and is worked on to a degree out of proportion to any likely use. While it may be too much to call it an “early work of art”, it is at least suggestive of an aesthetic sensibility. I suggest that we have here evidence for a very deep history of aesthetic production: a history so long that it makes the Upper Paleolithic look positively contemporary. This history extends back long before our own species emerged, long before language developed, long – apparently – before any genuinely symbolic activity of any kind.”

using creatures, we can't avoid the "cultural modulation" of our art behaviors. Soon enough, art will acquire many other functions and values; values dependent on representational properties, which bring with them a host of connections to religion and other social and political structures. The first musical sounds produced with bone flutes of the Paleolithic¹¹¹ would soon give way to culturally inflected ways of music making that accumulate layers of regulative conventions, bringing to bear centuries and millennia of experimentation on the unsuspecting improvisations of the illiterate musician that simply plays an instrument according to the musical tradition in which he grew; thus we can hear medieval plainchant in the practice of unsuspecting groups of rural singers; or think about the way the most bare and unsophisticated folk song of today depends on the conventions of equal temperament and diatonic tonality, which developed in the course of centuries (or how it may be a remnant of older practices). A vivid and more detailed example of this in music history was given by Alex Ross in his recent book *Listen to This*, second chapter, "Chacona, Lamento, Walking Blues", which traces the historical recurrence of a bass line, from the court music of renaissance Spain to modern day blues music.

Davies also points in the right direction when he argues that function (causal-agentive artifactuality) and procedure (status-functional artifactuality) can diverge as soon as an art practice is established. This opens an interesting possibility: as soon as a social practice of art making is established, we can have a status-functional concept of art that does not conflict with the concept of causal-agentive art, in the sense that the attribution of the status-function tracks, so to speak, the causal-agentive function (in Wollheim's terms, the presence of the causal-agentive function constitutes the *reason* for status-conferral). As Gregory Currie observed in an article on "Paleolithic art":

"the separation (if there is one) between the aesthetic and the artistic seems to be extremely recent and it can hardly be a criticism of any theory that it looks for connections between art and the aesthetic in the distant past, when virtually all but the last fifty years of art history reinforces that connection."¹¹²

I would further reinforce this idea by pointing out that the vast majority of art in the last fifty years, even in the western artworld, even within the so-called "avant-garde", does *not* suggest the disconnection between the aesthetic and the artistic. The idea that it does results from an egregious

¹¹¹ Bone artifacts capable of producing an array of discrete pitches date as far back as 30kya to 40kya.

¹¹² In S. Davies et al. (eds.) 2009, p. 2.

philosophical exaggeration of the (philosophical) importance of a minute portion of artworks produced in the twentieth century. It results from a particular form of conceptual illusion that seems stronger in theorization about the arts than anywhere else: the idea that present art is more important than past art, and that the “pre-theoretical intuitions” of avant-garde artists of the present, in the western artworld are more important than the pre-theoretical intuitions of the illiterate Yakutian shepherd playing the khomuz (a Jew’s harp). I take them to be on an equal standing. The painters at Lascaux, Altamira and Chauvet did not have the benefit of reading Danto, but then again, we don’t think less of Plato for not having read Kant.

There is a concept deployed by Jerrold Levinson in his historical definition of art that can be of use to us in explaining what is common to status-functional art that eschews art’s (primary) causal-agentive functionality and to art where this separation does not occur. This is the concept of a *regard as a work of art*; Levinson (1990, p. 9) uses this concept to define present art (art at *t*) as objects which “some person or persons non-passingly intended for *regard as a work of art*, i.e. a regard in any way (or ways) in which objects in the extension of ‘artwork’ prior to *t* are or were correctly (or standardly) regarded.” This view has been extensively criticized for allowing irrelevant properties to count in favor of something’s being art. For instance, if the Eiffel Tower is an artwork, then one of the ways in which it is regarded is as a tourist attraction. This seems to make the Eiffel Tower replica in Las Vegas to fit the definition (personally I have no objection to count this as an artwork, though a less than mediocre one; but then again, most artworks, like most artifacts of any kind, will not be outstanding ones). The clause about the regard intended being “correct” or “standard”, I gather, is meant to deal with such objections, though it seems that if we can give an account of what constitutes “correctness” here we begin to abandon the historical format of the definition. Now, quite apart from the discussion of the relative merits of the historical definition, the concept of *regard as a work of art* gives an important element of what may constitute a status-functional concept of art. The point of status-functions (so Searle has argued) is to generate deontic powers (positive and negative), i.e., rights and obligations. That something counts, say, as a territorial border may give us a desire-independent reason for behaving toward it in a certain way (e.g. refrain from crossing it in certain contexts). This reason is independent of our actual desires to cross it (or not to) and is also independent of whether there was a good reason in the first place to settle that particular territorial border (when the Berlin Wall existed, its existence gave people desire-independent reasons for behaving in a certain way towards that fact; this doesn’t mean that the existence of the Berlin Wall was a good thing). One initial counterexample to Searle’s conception of status-functions as a tool for generating powers was the case of purely honorific titles, i.e., titles that don’t seem to confer any positive or negative deontic powers. Think of the case

where someone wins a beauty contest or being awarded an honorific degree by a university. But we *can* think of a purely honorific title as carrying deontic power, to the extent that the person or thing to which the title was bestowed is to be *regarded* in a certain way (whether we actually observe this or have a good reason not to, is a distinct matter, like in the case of the Berlin Wall). This is the concept that seems to be at work in the status-functional artifactuality of artworks: we may think that a particular piece does not *deserve* the title of artwork; nevertheless, we don't behave towards that piece in the same way we behave towards just any artifact that we find worthless. We struggle to "get the point" the artist was making in presenting that object; we even consider immoral to destroy artworks merely on the grounds that we don't enjoy them. The more radical cases of art vandalism don't defeat this point, just as smuggling goods across a territorial border does not impugn the fact that the relevant status-function is at work, quite on the contrary: there would be no *transgression* if there was no status-function. We don't count it as a "big deal" if someone destroys a dysfunctional coffee-machine, but even art vandals consider their actions as constituting an important "statement"; this is only intelligible if we think some transgression was perpetrated.

In sum, when artwork status is conferred upon an object, some desire-independent reasons for acting are generated; in this sense, there is nothing inconsistent with the position that sees artwork status as a status-function and not merely an indicator that the causal-agentive function is present. Objects with the artwork status-function are to be *regarded* differently from other objects from which they may be physically or perceptually indistinguishable. The artwork status-function is thus of the *honorific* kind, it generates honorific deontic powers, or in other words, it generates a desire-independent reason to *regard as a work of art*, something that even the audience member most inimical to avant-garde experiments will observe, in thinking that his bashing comments make a relevant point. This is similar to the way that the actions of a political dissident attest to the presence of the system of status-functions to which he opposes. To recognize the existence of an institutional reality is not coincident with the ideological endorsement of it. Status-functions exist to the extent they are *recognized*, and recognition here is not synonymous with acceptance; recognition requires only that people *behave* as if the status-functions are in place.

At the same time, the concept of the artwork status-function as generating honorific deontic powers allows us to see how this concept can't take metaphysical precedence over the causal-agentive concept of art. In this way, the artwork status-function is different from the constitutive rules that specify the state of success in games (and other activities that are primarily status-functional). Namely, the causal-agentive concept acts here as the *source* of the status-functional concept; we get the status-functional concept by *abstracting from* the causal-agentive functions in virtue of which art first came into being. Yet, the causal-agentive functions form the background

against which the status-functional concept is intelligible as such. Art that is purely status-functional is *dependent* on art that is not purely status-functional. It is not that before the twentieth century art did not have an element of status-functionality. This element arises as soon as human beings regard aesthetic artifacts (whether or not these are *primarily* intended to have the relevant causal-agentive functions of art); the novelty of the twentieth century artworld was the conscious attempt at severing both elements.

At this point one question arises: *why* would we be interested in separating the status-functional from the causal-agentive in the domain of artworks? One answer to this question is that doing so allows artists to make interesting *artistic statements*. Art always was multivalent, that is, it always served a vast array of purposes: ideological, religious, etc., *as well* as purely aesthetic purposes. Art in the twentieth century was no different: it was used to bolster political ideologies as well as to ridicule them. One function that twentieth century artworks took was to question the very nature of art. This is the most noticeable achievement of some “purely status-functional” artworks, e.g., ready-mades. To make us think about the relations between the aesthetic and the artistic is another example of this. But there is also a surprising aspect lurking here: by forcing the artwork status-function on objects that seem to *lack* any trace of the causal-agentive functions typical of art, the artist may either consciously or unwittingly be generating *aesthetic properties* that those objects would otherwise lack. (One example could be that Duchamp’s *Fountain* is “witty”, while an identical urinal in a plumber’s shop is not.) This is a contentious point, since it depends on how wide we take the dependence base for aesthetic properties to be. If we think that aesthetic properties depend *solely* on perceptual properties, we’ll be inclined to consider that ready-mades cannot aesthetically differ from their non-artwork indiscernible counterparts; even if they differ in relevant *artistic* properties (e.g. ready-mades may have *meanings* that their non-artwork indiscernible counterparts lack – they are *about* something, in the way Danto and other philosophers took them to be). I leave this question open, since whatever our inclinations are in this matter, it will not impugn the thesis that the status-functional element in artworks crucially depends on their causal-agentive functionality.

In light of the above, I am sympathetic to Zangwill’s critique of Iseminger’s theory. The social practice of art (like all social practices) must be explained; otherwise the existence of art institutions is an unexplained explainer. While some social phenomena only emerge out of cooperative behavior and their nature is “wholly internal” to the application of constitutive status-functions

(they “emerge out of game-theoretic pressures on joint behavior”¹¹³), not all social practices are like this. Importantly, art is not like this. The social practice of art is not like games in the sense of sharing their “logical matrix” and the concept of causal-agentive (aesthetic) art takes metaphysical and explanatory precedence over status-functional art. I may disagree with Zangwill to the extent that I see no (absolute) need of some strategy to integrate recalcitrant cases into an aesthetic theory of art. But then again, to the extent that Zangwill is not primarily concerned with the “game of definition and counterexample”, the disagreement may be negligible.

One more relevant point about the distinction between the status-functional concept of art and the causal-agentive concept of art. Each of these concepts allows us to build two distinct, equally legitimate definitions of art: institutional art (or “artworld art”) and aesthetic-functional art. The latter subsumes other important functions that were regimented by previous “functional theories”, such as expression and representation, to the extent that our interest in such properties in art is typically an interest in the *aesthetically apt* realization of such properties. Suppose that we organize an exhibition that only considers applications of ideologically motivated artworks. Even within this restriction, a good deal of selection will have to be made; some works will be preferred in detriment of others. Since all of them satisfy the ideological criterion, the selection must be based on something else. This something else on which selection is based will plausibly be *aesthetic value*. The more aesthetically apt, ideologically motivated artworks of the lot will figure in the exhibition (this is compatible with the defeasibility of the jury’s verdicts: maybe some of the worse works will be selected and some of the better works rejected; but this does not impugn the fact that the operative criterion is an aesthetic one).

The set of institutional art (or “artworld art”) will also include many works that have the relevant causal-agentive functions (namely, aesthetic functions). Most artworks in art history will belong in the intersection of both groups – they were made to have the relevant causal-agentive functions *and* they have the honorific deontic powers that result from their status as artworks. Only a minute portion will be “recalcitrant” works that were made with a conscious effort to break “function” and “procedure” apart. But the crucial point is that each concept will include things that the other concept excludes. Aesthetic functionalism is often criticized for *including too much*; namely, cake decoration, tattooing and fireworks, among other examples, will be subsumed under the aesthetic-functional concept but excluded from the status-functional concept. This is a perfectly acceptable state of affairs, given 1) art concept pluralism (we have more than one consistent concept of art) and

¹¹³ Zangwill (2009: pp. 39-40).

2) the metaphysical precedence of the causal-agentive concept over the status-functional concept. This is consistent with an important historical and sociological fact: in different times and places, different institutional groupings of the “arts” (the practices of producing aesthetic artifacts) were consolidated. The concept of “Fine Arts” is an eighteenth century invention, namely coming into being with works such as Charles Batteaux’s book of 1746, *Les Beaux-Arts Réduits à Un Même Principe*.¹¹⁴ From ancient Greece to the Renaissance, *music*, for instance, was grouped with mathematics and both the Greek word *techne* and the roman word *ars* were applied to practices that are excluded from the eighteenth century concept. In traditional Japan, things like calligraphy, “rock gardens”, sword making, martial arts, flower arrangements, incense and the tea ceremony, along with other practices more akin to “our” concept of *art*, like woodblock printing and painting, have an importance (an element of status-functionality) that is absent from the successive western groupings of the arts.¹¹⁵ The distinction between the causal-agentive concept of art and the status-functional concept of art allows us to make sense of these contrasts, without ending up with the conclusion that “art” is not a consistent concept, or that there is no operative definition or conception of art other than an incorrigibly vague notion based on “family resemblances”. We can have illuminating distinction and explanation rather than confusion and fuzziness.

Considerable shifts in how the several practices of *aesthetic craft* (“art” in the causal-agentive sense) are grouped, and how they are conceived in terms of status-functionality (cross-cultural nuances in the concept of *regard as a work of art*), is precisely what is to be expected given a) the metaphysical primacy of the aesthetic-functionalist (causal-agentive) concept of art over the status-functional concept (which gives us the result that there are always potentially much more kinds of *aesthetic craft* than can be “encoded” in institutional groupings of “the arts”); b) art’s multivalence (aesthetic functions can be “modulated” by an immense variety of non-aesthetic functions – religious, ideological, philosophical, etc.); c) the vast array of values, other than aesthetic value, that can enter into mutual dependency relations with aesthetic value.

Another welcomed result of this discussion is that cross-cultural differences in the possession of an explicit *art concept* do not pose any substantial challenge to a cross-culturally applicable theory of art. The existence of many “art games” or *institutional groupings* of the more basic practices of “aesthetic craft” are not only consistent with the metaphysically more basic aesthetic-functionalist concept but *require* it to be intelligible as social practices.

¹¹⁴ See Paul Kristeller (1965: p. 199): “The decisive step toward a system of the fine arts was taken by the Abbé Batteux, he was the first to set forth a clear-cut system of the fine arts in a treatise dedicated exclusively to this subject.”

¹¹⁵ See Zangwill (2007: pp. 75-76).

2.3 Descriptivism, revisionism and the pragmatic constraint

According to a methodological view that has come to be known in the literature as the “pragmatic constraint” on the metaphysics of art kinds, philosophers should avoid writing, for instance, “of music as if it possesses a universal and unchanging essence (...) as if what we believe about musical works, how we respond to them, and how our social practices and histories incorporate them are irrelevant to the kind of things they are, just as a ruby is what it is even if humans believe it to be the frozen blood of the gods.”¹¹⁶ Succinctly, it is the view that philosophical speculation about the arts should be constrained by the actual practices of art, that is, philosophers should bear in mind that they are theorizing about a social practice, not a natural kind whose boundaries, whatever they may be, are quite independent of what human agents think and do.

The pragmatic constraint is connected to a wider methodological thesis about the philosophy of art, i.e., metaphysical descriptivism. Even though closely related, the two should be distinguished. Descriptivism can be taken either as a general view about the metaphysics of social and artifactual kinds (which would then apply to art and artworks in virtue of these falling under the former), or as a special view concerning exclusively the nature of art and the ontology of artworks (this view can be termed “local descriptivism”). Descriptivism is opposed to metaphysical revisionism; what Julian Dodd, one of its defenders, calls “meta-ontological realism”, or Amie Thomasson, one of its foes, “the discovery view”. Though these two methodological theses are most often discussed apropos the subject of musical ontology (and artwork ontology in general), and less so concerning the topic of definition, they are clearly equally relevant in both cases. For simplicity, I'll use the term “descriptivism” to refer to art-metaphysical descriptivism, encompassing both ontological and definitional issues. Succinctly stated, descriptivism is the view that our metaphysical theories about the nature of art should be constrained by the pre-philosophical beliefs about the nature of artworks, either explicitly held or latent in our critical and appreciative practices. The relevant pre-philosophical beliefs can be identified either with “folk metaphysics” – the beliefs that most people have on the subject – or with the “spontaneous metaphysics” of artists, professional critics and similar authorities – the authorities that most people would be inclined to defer to in their deployment of art related concepts. In other words, descriptivism is the view that the conclusions

¹¹⁶ S. Davies (2008: p. 373). This quote was originally directed at Julian Dodd in a discussion of timbral sonicism.

one can legitimately infer via armchair philosophical speculation about the nature of art must be bound by the beliefs of those who “ground and reground reference” of our art kind terms.¹¹⁷

The motivation for metaphysical descriptivism is a disputable idea based on a plausible observation. The latter is that social and artifactual kinds are relevantly different from natural kinds; the former is that the ontological nature of these kinds is determined by the beliefs of human agents. Basically, it is the idea that we can’t make significant metaphysical discoveries about social and artifactual kinds that aren’t already “contained”, explicitly or implicitly, in our pre-theoretical beliefs about them. In other words, while it may turn out that some of our terms for social and artifactual kinds fail to have reference, i.e., that they fail to pick out any items in the real world, it can’t turn out that we are mistaken (or radically mistaken) about their nature, i.e., about what they are, should they really exist. This is neatly illustrated by the example of dreamcatchers, used by Julian Dodd in an argument against metaphysical descriptivism.¹¹⁸ A dreamcatcher is an artifact that has the function of catching bad dreams. As it happens, the world is such that no artifact can fulfill the function of catching bad dreams, so that, in a sense, there are no things that are dreamcatchers; the artifactual kind term “dreamcatcher” fails to pick out something that actually catches bad dreams. Yet, in another sense, there are dreamcatchers: there are artifacts which we refer to as “dreamcatchers” and about which some people (supposedly) believe they have the power of catching bad dreams. It is arguable that an artifact’s inability to perform its intended function does not prevent it from *having* that function. This is a plain consequence of the *normativity* of artifactual kinds: artifacts can be *defective* without thereby losing their identity. Nevertheless, one could make an argument from the *impossibility* of anything at all performing a given function to the inappropriateness of the intentions to impose that function on any object in the first place. After all, if I intend to make a time machine with a few pieces of wood and metal from a junkyard, that will not suffice for the product of my efforts to be appropriately named a time machine. A response to this argument is that since people can appreciate dreamcatchers for qualities other than the ability to catch bad dreams, there need be no general belief that dreamcatchers catch dreams in order for them to be made and used, which contrasts with the example of my ludicrous and utterly useless time

¹¹⁷ Thomasson: (2005: p. 226): “The ontological status of paintings, symphonies, or works of literature is thus not something we can discover by investigations into the mind-independent world; instead, we must follow the method of analyzing the conception embodied in the practices of those competent speakers who ground and reground reference of the term.”

¹¹⁸ Dodd (2012): “Consider a possible world in which we all believe in the efficacy of dreamcatchers: that is, those nets decorated with beads and feathers that the Ojibwa people supposedly believed to prevent bad dreams. In such a world, dreamcatchers are produced in order to catch bad dreams: this is the purpose with which they are produced and which the wider public takes them to perform.”

machine. Another response lies in an analogy with religious artifacts, such as holy water. It is not necessary to believe that holy water has any supernatural powers in order to believe that some water is holy water; after all, it is true of some water that it has been blessed by a priest, and such water is then appropriate for religious use, whether or not it possesses supernatural properties in virtue of possessing that one relational property. This example complicates matters, since some artifacts have simple artifactual functions that require only a link to appropriate intentions and an adequate physical structure to perform the intended function; but other artifacts are made to perform institutional functions, functions whose performance requires convergence in the representation of function. This seems to be the case with holy water. It is an artifact with a procedural ontology rather than a simple causal-agentive ontology. Let's drop that example for now.

The point of the dreamcatcher example is that whether or not the artifactual kind term "dreamcatcher" successfully picks out an item in the real world, the agent who intentionally builds a dreamcatcher cannot be wrong about what a dreamcatcher is. The world may fail to fit the agent's creative insight or "mental template" for the artifact in question, in that nothing that issues from the agent's efforts will have the causal efficacy required to perform the intended function, but the agent cannot be wrong about the "mental template" itself, i.e., about *the kind of thing that it specifies*. Where all this aims is that we can't make significant discoveries about the ontological nature of the referents of our artifactual kind terms, in the same way that we make significant discoveries about the referents of our natural kind terms. We may discover that water has a certain chemical structure, that silver is a natural element with atomic number 47, or that whales are a kind of mammal, not a fish, but similar discoveries are barred to us in the domain of artifactual kind terms. What this means is that we have a sort of *epistemic privilege* about the latter that we lack about the former. The controversial core of this suggestion is that the "mental template" just mentioned contains a "nascent ontological conception" of the kind in question, and that the metaphysician's work consists in reading off from this "nascent conception" the constraints that must characterize the kind in question (hence the qualification "descriptivist"), so that, if the "mental template" determines that, say, musical works are repeatable, creatable, and modally flexible entities, no philosophical enquiry should exclude one of these characteristics on purely metaphysical grounds. In other words, it can't turn out that musical works are repeatable, non-creatable, modally inflexible entities. There are no significant empirical or metaphysical discoveries to be made in the case of social and artifactual kinds that could lead us to revise our "mental template", in the way that there are such discoveries to be made in the case of natural kinds, that can lead us to revise our "mental template" of them, e.g., that whales are really a kind of mammal, not fish. For the metaphysical descriptivist, social and artifactual kinds can't turn out to be at variance from our "mental templates" for them since, unlike

natural kinds, social and artifactual kinds would not exist if not for our “mental templates”: social and artifactual reality is the making of human agents; it is causally downstream from our “mental templates”.

We can illustrate the descriptivist’s point by appealing to concepts which are “culturally inflected”, even when they are not social or artifact kind concepts: we could use the concept mountain differently, but we can’t “discover” that things that did not previously count as mountains now, surprisingly, do count. What we can do is change our criteria for counting mountains. But since mountains, in Stephen Davies’ terms, “latch on to” natural, mind-independent structures, there are discoveries we can make about them once the criteria for using the sortal “mountain” are fixed: e.g. we can discover that they are the result of large scale geological processes, such as tectonic drift, volcanic activity and erosion.

Social and artifactual kinds are twice removed, so to speak, from “natural structures”, so that our “mental templates” for such entities are determinant in ontological terms, and not merely in terms of how we divide a previously existing reality, since in this case there is no previously existing reality to divide. Another example is the case of chess: the rules of chess are constitutive conventions, which means that our “mental template” for chess does not “carve up” a previously existing reality; rather, it makes it the case that there is such a thing as chess. If we want to know what chess is, we must produce an accurate description of the template. Any metaphysical “revision” of the template will only change the subject.

The “discovery model” (also called “meta-ontological realism” or “metaphysical revisionism”), on the other hand, is the view that the dependence relation between social and artifactual kinds and human agents’ intentions or convergence in representation affords us no such epistemic privilege. We can bring out the revisionist’s point via an analogy with beliefs: we may have privileged access to the content of our beliefs, in the sense that it is transparent to me that my belief that Mount Fuji is located in Japan and is over twelve thousand feet high is a belief about a mountain, not about, say, an array of lines and colors in a woodblock print. But this fact has no bearing on my knowledge of what beliefs are in themselves. There is room for widely revisionist conceptions of what a belief is, or to what ontological category my occurrent belief that Mount Fuji is located in Japan and is over twelve thousand feet high belongs. Our beliefs about beliefs do not fix the ontology of beliefs. This may not be a perfect example, since the concept *belief* is not a social or artifactual kind concept, but the point can be easily extrapolated, says the revisionist, to the case of such concepts. Take the case of symphonies: the use of the concept suffered significant changes from Bach’s time to the twentieth century. The sense in which the opening section of the *C minor Partita* for keyboard, BWV 826 is called a “symphony” is a world apart from the sense in which Beethoven’s symphonies

are thus termed. Our use of the concept is only fully intelligible against the historical background of the development of musical form, of which a crucial component is the concept of sonata-form. What counts as a symphony crucially depends on historical relative conceptions of musical form. Though the conventions regarding symphonic form may change in drastic ways, we can't *discover* that something not previously counted as a symphony turns out, surprisingly, to be an example of the form. Our "mental template" deployed in the conventions about musical form determines what musical entities count as symphonies; there are no symphonies independently of such mental or conceptual template. Yet, how can this fact allow us to infer that those things that count as symphonies according to the convention are correctly described as abstract, creatable, modally flexible entities, rather than abstract, non-creatable, modally inflexible entities, or, for that matter, as a set or mereological fusion of concrete entities, or any other ontological characterization? Why should there be no room for error about these features just because there is no room for error on what counts as a symphony (except in the fairly innocuous sense that someone ignorant of the musical conventions may confuse a symphony with another "template")? How can the constitutive conventions of chess help us in infallibly placing the game of chess itself in one ontological category rather than other?

The main descriptivist argument for such conclusion was presented by Amie Thomasson¹¹⁹ as a consequence of solving a problem in the disambiguation of reference that arises for causal theories of reference. This problem is known as the "*qua* problem"¹²⁰. According to causal theories, the reference of both singular and general terms that function as names is not determined by an associated descriptive content such as "the teacher of Alexander" for "Aristotle", but by a causal chain that traces back to an originating use of the name or "baptism" of the thing named, where the name is ascribed to an object by way of ostension, e.g., "We hereby name you Aristotle". According to the causal theory, when we use the name "Aristotle" we are deferring to the way previous users employed the name, going back to the originating use, and it is in virtue of this causal chain of uses that the name successfully picks out the ancient Greek philosopher who was the teacher of Alexander. Likewise, those who ground the reference of natural kind terms like "water" do it by an act of ostension in what is the originating use of the concept. Further, non-originating uses of the term involve deference to previous uses. The term "water" refers to water in virtue of standing in such a chain of non-originating uses that ultimately trace back to an originating use, where the reference of the term was grounded by ostension: pointing to a sample of water, someone

¹¹⁹ See Thomasson 2005. See also her "Artifacts and Human Concepts", in Margolis, E. & Laurence, S. (eds.) (2007).

¹²⁰ See Devitt & Sterelny (1999: pp. 79-81; 90-93).

utters “We hereby name *that kind of thing*, ‘water’.” (Evidently, this is a rational reconstruction, not an accurate description of how the term was *actually* introduced; the point is that the *mechanism* at work was of the sort described.) This picture of how a term’s reference is grounded allows for substantial error, on the part of those who “ground and reground” the term’s reference (i.e. those who originate the use and those who keep non-originating uses in tandem with the originating use), concerning the nature of the kind of thing referred, and thus it allows for significant discoveries about the kind in question. For instance, we discovered that the term “water” refers to that which is H₂O and that it did so even *before* we knew about it; we discovered that the term “whale” refers to a kind of mammal and that it did so even when people thought that whales were a kind of fish. To recall Stephen Davies’ example from a while back: even if human beings believe rubies to be the frozen blood of the gods, the term “ruby” (and any word expressing the same concept) still picks out a crystalline form of aluminum oxide, whether or not users of the word ever become aware of the fact. In sum, widespread error of this kind does not constitute an impediment to a term with the associated erroneous information successfully referring to its target.

Now, the *qua* problem arises when the context of originating use allows for ambiguity regarding what, if anything, the causal chain traces back to; a problem that purportedly arises in its “full virulence” in the case of social and artifactual kinds.¹²¹ What this means is that in any context of originating use, the would-be grounders of a term’s reference will be faced with an indefinite number of instantiations of different kinds of things, all of which are equally good candidates to stand at the originating end of the causal chain of reference. Robert Stecker illustrates the *qua* problem with a neat example:

“At the original name-giving occasion, or, for that matter, on some other occasion in which the name is used, how does it get hooked up with the right object, since there will always be other objects present? Suppose I point to my newly acquired favorite rock, and I say, ‘I name you Rocky.’ I have pointed in the direction of any number of other things not identical to the rock: a collection of molecules (Rocky could lose some molecules and still be Rocky), a rock surface, and so on.”¹²²

And a straightforward generalization of the problem is given by Amie Thomasson:

¹²¹ Thomasson (2005: p. 222).

¹²² Stecker (2009: p. 380).

“whenever a grounder attempts to establish the reference of a term, there will in fact be a great many ontological kinds of thing available for reference, from biological individuals, to collections of particles, to abstract types, created types, qualities, events, and so on.”¹²³

In other words, the attempt at grounding reference via ostension will be ambiguous between an indefinite number of equally legitimate “ontological candidates” to stand as the source or origin of the causal chain of reference. One way of disambiguating reference is by appealing to a minimal descriptive content associated with the term. By naming his favorite rock “Rocky”, the speaker must have in mind a “background ontological sortal”,¹²⁴ without which his attempt to refer anything at all will be unintelligible as such. To count things in any given context, one needs sortal concepts; the question “how many *things* are here” is unanswerable, unless we are implicitly operating with a set of background sortals (e.g. the question is understood as being about perceivable, medium-sized concrete particulars that can be easily moved around, and so on). An illustration of this is found in Putnam’s thought experiment of “a world with three individuals”, which is described differently by two persons operating with distinct background sortals, even though they express these different sortals using the same word: “object”. While the first asserts that there are three objects in the imaginary world (x_1, x_2, x_3), the latter asserts that the said world contains seven objects, since he believes “that for every two particulars there is an object which is their sum”¹²⁵ so that besides x_1, x_2 , and x_3 we also have $x_1+x_2, x_1+x_3, x_2+x_3$, and $x_1+x_2+x_3$. The apparent incompatibility of both descriptions results from the fact that different background sortals are being used: one according to which sums of objects count as objects, another according to which they don’t. It doesn’t follow from this that the descriptions are incompatible (genuinely incompatible descriptions use the *same* sortals) nor that there is no fact of the matter as to how many objects are there in a given context. What happens is that the question “How many *things/objects* are here?” is unanswerable until we associate “thing” or “object” with a background sortal. Most often we implicitly associate the word “object” with a background sortal for concrete particulars of a certain kind (things with clear spatial boundaries, which are not events). Recall once again the sortal term “mountain”, for which we saw there are no clear-cut conditions of individuation. Consider a case where the same mass of rock is configured in such a way that it is not entirely clear whether it counts as one or two mountains. In that situation, two persons could associate different constraints with the count noun “mountain”, so

¹²³ Thomasson, *op. cit.* p. 226.

¹²⁴ Kania in Stock & Thomson-Jones (2008: p. 30).

¹²⁵ Putnam (1992: pp. 96-97).

that one of them would count one mountain and the other two. What they produce are not two incompatible descriptions, since arguably they are using two distinct but homonymous sortals, equally satisfied by the same chunk of matter.

Now, imagine that instead of naming his favorite rock, as in the previous example, our would-be grounder is naming his last artistic achievement (not that Rocky wouldn't make a fine ready-made), say, a string quartet:

“For instance, when Shmarb proclaims ‘*This* is my sixth string quartet’, pointing to the pile of inked-up manuscript on his desk, if his baptism is to succeed, he must be thinking of a whole, multiply instantiable, sonic item – as opposed to a concrete particular (to which he is also ostensibly pointing), a spatial or temporal part of that concrete particular, the countersubject of the fugue in the fourth movement, and so on.”¹²⁶

In other words, by naming one of the items that equally satisfy his act of ostension, Shmarb must have a background sortal if reference is not to be indeterminate. And in naming a string quartet, he is deploying a host of conventional concepts rooted in a social practice of music making. It will not be sufficient, in order to disambiguate reference, to specify that it is an art kind he is referring to, since the context will provide an equally gerrymandered array of candidates: he could be referring to a sound event type by pointing to a representation of it (the score), to an action type, and so on. To disambiguate reference, Shmarb's use of the sortal term “string quartet” must have associated with it a minimal descriptive content derived from the history of musical practice, a “nascent ontological conception” that singles out, among all possible candidates, the art kind he refers to. And what goes for string quartets also goes for paintings, symphonies and any other art kind.

“The disambiguating conception of the sort of kind involved in grounding the reference of art-kind terms must include a nascent concept of the ontological status of the kinds of work involved. As a result, at least a background concept of the ontology of the work of art is needed to establish the reference of terms like ‘painting’ or ‘symphony.’ Such concepts determine the ontological kind, if any, picked out by the term, and so the ontology of the work of art must be something we learn about through conceptual analysis of the associated concepts of people who competently ground (and

¹²⁶ Kania, *ibid.*

reground) the reference of terms like 'symphony' and 'painting,' not something we can seek to discover through investigations into mind-independent reality.”¹²⁷

What this amounts to is that the solution to the *qua* problem requires an hybrid theory of reference, according to which reference is fixed not merely by an external causal relation to the context of originating use via a chain of deferring non-originating uses, but also by a minimal descriptive content (a “background ontological sortal”) without which a causal chain of reference cannot be initiated.

The core thesis in this story for art-metaphysical descriptivism, needless to say, is the idea that the background ontological conception that grounders of reference associate with art-kind terms *determines* the ontological kind picked out by the term, if it refers at all. The idea is that we don’t first refer to a certain kind of thing and then proceed to discover mind-independent ontological facts about it (in the way that we proceed to make empirical discoveries about whales), since reference is only successful in the first place in virtue of the “conceptual template” or background sortal that guides the act of ostension. The idea here is that we must already *know*, if only implicitly, what we are referring to, so that setting out to make ontological discoveries about an art-kind doesn’t seem like a consistent endeavor to engage in. Supposedly, such project would be almost tantamount to trying to discover how many angles a triangle has. The knowledge is already built in our ability to refer to triangles. The analogy is not perfect, though, since the status of assertions like “paintings are unique concrete particulars that can be forged” or “musical works are non-eternal abstract continuants with multiple instances” differs from the status of assertions like “triangular figures have three angles”, in that the latter is true by definition while the former two are not. The assertions about paintings and musical works above do not present a semantic equivalence, that is, “unique concrete particular that can be forged” is not the *meaning* of “painting”, in the way that “having three (an only three) angles” is what “being triangular” means.

The metaphysical descriptivist is not engaged in providing lexical definitions for art-kind terms but rather in working out a description of the properties of art kinds themselves, which most accords with the way those kinds are conceived in our artistic and appreciative practices. These are taken to constrain the metaphysics of art, and the motivation for such view is that we can’t work out the “basic ontology” of art kinds independently of some *epistemological* questions about art, in particular questions about what we take to be the main focus of appreciation in artworks. In other words, the view that underlies metaphysical descriptivism is that the epistemology of art kinds

¹²⁷ Thomasson (2005: p. 223).

fixes, to some extent, the metaphysical facts about art kinds. This is the crucial thesis underlying the “pragmatic constraint” on the philosophy of art kinds.

An explicit articulation of the pragmatic constraint was given by David Davies in his 2004 book, *Art as Performance*:

“Artworks must be entities that can bear the sorts of properties rightly ascribed to what are termed ‘works’ in our reflective critical and appreciative practice; that are individuated in the way such ‘works’ are or would be individuated, and that have the modal properties that are reasonably ascribed to ‘works’ in that practice.”¹²⁸

Or more succinctly:

“Artworks must be conceived ontologically in such a way as to accord with those features of our critical and appreciative practice upheld on rational reflection.”¹²⁹

Both David Davies’ formulation of the pragmatic constraint and Thomasson’s argument from the *qua* problem are particular applications of the more basic view that the epistemological facts fix the metaphysical facts about art kinds. This feature is what (supposedly) is most markedly distinctive of social and artifactual kinds as opposed to natural kinds, whose “basic ontology” can be settled independently of epistemological considerations. Both Davies and Thomasson acknowledge that a minimum degree of revision may be required to make our descriptions of art kinds consistent – thus Davies’ clause about the features in terms of which we conceive of works ontologically being “upheld on rational reflection”. This already raises the problem of how much is to be decided by “rational reflection” before the pragmatic constraint is violated and a “change of subject” takes place.

The role of art metaphysics, for the descriptivist, is not to discover surprising “ontological facts” independent of how we think about art kinds, but to describe the metaphysical constraints already built in the “conceptual templates” of those competent speakers who ground and reground the reference of art-kind terms. Any ontological revisionist move here would amount to changing the background sortal associated with the term, which would then most likely change the referent and thus the subject of investigation. If examining art-kind terms as they are used in our art-historical

¹²⁸ D. Davies (2004: p. 18).

¹²⁹ D. Davies (2004: p. 23).

practices yields the conclusion that the “background ontological sortal” that competent users associate with the term “musical work” is “non-eternal, creatable abstract continuant” we cannot subsequently discover that musical works are really eternal, modally inflexible abstract types.

This core thesis can be challenged, though. One way to do this would be to challenge the need for a hybrid theory of reference. But this has the disadvantage of committing the metaphysical revisionist with a “pure” causal theory of reference for social and artifactual kinds. A stronger solution would be to show the compatibility between the need for disambiguating descriptive content and metaphysical revisionism. This is in fact the road taken by Julian Dodd, one of the staunchest defenders of metaphysical revisionism. He argues that Thomasson’s hybrid theory of reference isn’t the only one available, and that other alternatives allow us simultaneously to solve the *qua* problem and to preserve meta-ontological realism.

Even granting that an hybrid theory of reference is the correct solution to the *qua* problem, and that background ontological sortals are required to disambiguate reference, it is not clear that our “mental templates” for art kinds cannot contain errors or that none of the information we must associate with art kinds in order to make reference to them possible is subject to revision.

Notably, Julian Dodd (2012) has recently argued against Thomasson’s position on the grounds that her argument from the *qua* problem actually consists in two related semantic theses: 1) that the solution to the *qua* problem requires a hybrid theory of reference; 2) that the best hybrid theory available is a theory of reference by fit. He counters her second semantic thesis by arguing that Gareth Evans¹³⁰ hybrid theory of reference in terms of a *dominant source of information* associated with the relevant terms, is not a theory of reference by fit between a thing and our “mental template” (or “nascent ontological conception”) for that thing is compatible with the first thesis and it gives the meta-ontological realist comfortable dwelling space. Thomasson’s theory is more in line with Searlean internalism, in that it takes the information in one’s head to *determine* reference by fit: either the world satisfies the “nascent ontological conceptions” speakers associate with art-kind terms or the term does not refer at all. Her response to cases of reference failure then resembles Weitz’s theory of the open-endedness of “empirical concepts”: in such cases, we will either conclude that there are no such things as specified by our “nascent ontological conceptions” or a *decision* will be made to change the meaning of the term, so that art-kind terms will then refer to the “closest ontological competitor”. Suppose we discover that those creatures we believed to be a kind

¹³⁰ The reference here is the paper by Evans and Altham (1973), “The Causal Theory of Names”.

of animal – kangaroos – were in fact sophisticated Martian robots¹³¹. In Thomasson’s view, since the information associated with the term “kangaroo” has proved to be mistaken (the term would not pick out biological entities, and these have identity conditions quite distinct from those of artificial machines), the conclusion is that there are no such things as kangaroos and therefore our term “kangaroo” does not refer. If we decide to keep calling the Martian robots “kangaroos”, we will have settled for the “closest ontological competitor” (a kind of robot).

According to Evans’ hybrid theory of reference, on the contrary, reference is secured by the fact that *those things* that we saw hopping across hills and mistakenly took to be a kind of animal, are the *dominant causal source* of the (mistaken) information we associated with the term “kangaroo”. There is a threshold below which reference fails, but that still allows for substantial error while preserving reference. Of course, upon discovering that the hopping creatures are in fact deceptively sophisticated robots, people could decide to alter their use of the term, and there is no way to predict what they would do. But still, the question remains whether the term “kangaroo” *as was used heretofore* actually referred to those things people saw hopping across hills or didn’t refer at all. *That* fact can’t be altered by what people decide to do next with the term “kangaroo”. Take a related case: when Gell-Mann stipulated that the word “quark” refers to a kind of physical particle, he picked the word from James Joyce’s *Finnegan’s Wake*. But the act of stipulation is powerless to determine anything about the previous use of the word. Likewise, whether the term “kangaroo” actually referred to the hopping robots is fixed by whatever theory of reference happens to be true, not by what people later on decide to do with the word. If we take something like Thomasson’s theory of reference by fit to be correct, we will conclude that the term did not refer; if we take Evans’ theory of reference, we will conclude that the term referred to the hopping robots, *despite* the information associated with the sortal “kangaroo” being mistaken.

I’ve mentioned a threshold of error, below which reference would not be secured. It is unlikely that such a threshold can be specified *a priori*, since that will probably vary with each case. In the counterfactual case of the hopping robots, despite we mistakenly take them to be kinds of biological organisms, enough information is associated with the term to still secure reference. Despite being wrong about their nature as biological organisms, speakers would still associate many concepts which would be true of the things referred: *something that hops across hills and looks and behaves like an animal*, for instance, but there will be many more such concepts. Likewise, in the case of art-kind terms, there is much more information associated with those terms than the information that

¹³¹ The example is from Richard Miller (1992).

fully commits us with specific existence and identity conditions. True, we think of musical works as things that are created by their composers and did not exist previous to the compositional act or “indication” by the composer. But even if that information proved wrong, there is much more information associated with the sortal “musical work” that could secure reference despite our mistaken ontological assumptions: *something performable with musical instruments; something audible; something that can be heard at the same time in different places; something that was composed (represented in a score) at a given time; something that is for listening*; and so on and so forth. The most demanding and definite our “nascent ontological conceptions” will be, the more likely it is for them to contain mistakes. But here we must note that philosophers are prone to exaggerate the role of *intuitions*, since what is “intuitive” radically alters with context. When one starts learning an instrument, things that are so intuitive for trained musicians that become basic actions for them will be extremely unintuitive for the beginner, who stumbles continuously because he must consciously think simultaneously in many different things that the trained musician no longer has to consciously think through while playing (e.g. positioning one’s fingers in piano or guitar playing). The same goes for philosophy or any theoretical practice. After some time struggling with conceptual questions and working out conceptual distinctions, we find that our “intuitions” no longer match the “intuitions” we previously had. Some of these may turn out to be highly *counter-intuitive* after a while. The context can radically alter what is “intuitive”, a point often made by those who are keen to point out that “language constrains thought”, hoping to draw metaphysically momentous consequences from this trivial fact (of course our language skills constrain our ability to make fine-grained conceptual distinctions; otherwise we would not have to be *taught* the difference between, e.g., a necessary and a sufficient condition, or the difference between *if...then* and *if and only if*, etc.; neither would we be as prone to incur in fallacious reasoning, something that often happens even when one is able to draw fine-grained conceptual distinctions and consciously tries to *avoid* such mistakes). Furthermore, “intuitions”, of all the kinds of belief one can entertain, are eminently *defeasible*. Intuitions should always be taken “with a grain of salt”, given their malleability and context-sensitivity. Consistency with our pre-theoretical intuitions should not be considered an absolute epistemic virtue, to the point that we exclude a metaphysical theory based on the fact that it doesn’t fit our “intuitions”.

The point of this discussion of the argument from the *qua* problem, of course, is not to prove beyond doubt the truth of a hybrid theory of reference in terms of a dominant causal origin of our “nascent ontological conceptions” (the disambiguating information associated with art-kind terms). The point is that it is far from obvious that the *only* viable theory of reference is a theory of reference by fit. And there is no pragmatic constraint that allows us to evaluate theories of reference

bypassing philosophical discussion. But then the pragmatic constraint can't be used to exclude theories of art kinds, since the viability of the pragmatic constraint is itself dependent on issues that cannot be solved by appeal to a further pragmatic constraint – but perhaps some *other* plausible theory that explains how our “intuitions” about social and artifactual kinds determine their ontological nature could be substituted for the language argument as a base for a pragmatic constraint; in which case, of course, we would have to *philosophically* evaluate such theory, which kind of yields the same result – there is no “easy road” to art ontology, even if the right conclusion to be arrived at is that art ontology is a misguided endeavor. Unless of course the pragmatic constraint one has in mind is something so general that *no one* actually disputes it. Something like: “a theory of musical works must be *about* those things with which musicians, composers and critics are actually engaged.”

2.4 Should we replace the pragmatic constraint?

But what would be the point of doing philosophy if we were not open to the thought that we have hitherto been wrong about something? We don't have to think of something as what it is in order to think of it and talk about it.

Mark Sainsbury¹³²

Why should we accept the pragmatic constraint on art-kind ontology? More importantly: why should we accept the pragmatic constraint when it appears that philosophers who accept the pragmatic constraint still arrive at radically different ontological proposals about their subject matter?¹³³ Perhaps this is the wrong way to phrase the question. Endorsing the pragmatic constraint will not automatically give us the ontology of an art-kind; it merely provides us with a way to evaluate ontological proposals. Ontological proposals will vary in plausibility according to how much of the relevant practices they manage to preserve. But as I see it this idea seems riddled with problems from the outset.

First, there is the problem of what *version* of the pragmatic constraint one is to endorse. What *are* the “relevant practices” that should constrain metaphysical theorizing? Robert Stecker has identified two main divisions among philosophers who endorse the pragmatic constraint. The first group gives special emphasis to our language or *referential* practices. The best representative of this group would be Amie Thomasson, with her argument from the reference of art-kind terms against metaphysical revisionism. The second group is represented by David Davies, who emphasizes the interpretive and appreciative practices of “experts” such as artists and art critics.

It is not clear to me that we have even a clean division between two groups here. Whatever the arguments against metaphysical revisionism, defenders of the pragmatic constraint are committed with the view that there are some pre-philosophical beliefs about art kinds that are not susceptible of revision or at least not susceptible of *radical* revision, even if some degree of revision is justified. (This already poses a second problem: *how much* revision would be justified and what counts as justification for revision?) For what the language argument for metaphysical descriptivism aims to establish is that a solution to the *qua* problem makes it the case that common-sense beliefs about

¹³² Sainsbury (2009: p. 98).

¹³³ For instance, philosophers who endorse the pragmatic constraint have defended ontological views about musical works in which these are a) indicated types (Levinson), b) generative performances (D. Davies), c) historical individuals (G. Rohrbaugh).

art-kind terms (the “nascent ontological conceptions” associated with terms for art kinds) are either correct in their general outline and, to that extent, insusceptible of revision, or our art-kind terms do not refer at all, in which case, according to Andrew Kania, we should take fictionalism about musical works (the idea that discourse about musical works is not aimed at truth) as a serious alternative. Yet another possibility, one endorsed by Amie Thomasson, is that if there are no things with the properties ascribed by the “nascent ontological conceptions” associated with art-kind terms, these terms refer to the closest ontological competitor instead. This raises another host of problems: how *far* can the “closest ontological competitor” be until we say that an art-kind term has no reference? What if there are two or more (seemingly) equidistant ontological competitors?

Since in most uses of art-kind terms successful reference is assured by deference to the grounders of the terms (what makes it the case that someone unfamiliar with the history of western music refers successfully to a certain kind of musical form – a sonata for orchestra – when using the term “symphony”, instead of referring to, say, a kind of medieval instrument used to produce a drone accompaniment¹³⁴, is that her use of the term defers to other uses, tracing back to the originating use of the term) the view that emphasizes common-sense beliefs about art kinds seems to collapse into the view that it is the pre-philosophical beliefs of “experts” that really count, for everyone else’s use of art-kind terms defers to the use of the term by “experts”. Either way, we end up with a view according to which the “nascent ontological conceptions” of grounders and regrounders of art-kind terms somehow determine the ontological nature of the kinds referred to, if anything is referred to at all. The difference seems to be whether one endorses this view in virtue of the argument from language (descriptivism as the solution to the *qua* problem) or in virtue of some other argument for the ontological determination of art kinds by the pre-philosophical beliefs of those who ground the relevant terms. What we have in both cases is a theory of the *social construction* of art kinds: because art kinds are socially constructed, no substantive metaphysical *discoveries* are possible about art kinds. We come to know the ontological nature of art kinds by analyzing the relevant beliefs inherent in the practices that give rise to art-kind talk. The merit of the language argument (the argument from the *qua* problem) is that it specifies a mechanism by which this could be so, one that we can discuss and devise strategies to refute. The argument cannot be that descriptivism is true *given* that art kinds are socially constructed, since this is question-begging:

¹³⁴ “Symphony” is also the name of a medieval instrument similar to the hurdy-gurdy. The term “sinfonia” was also used before the eighteenth century to refer an introductory instrumental piece (see, e.g. the “Sinfonia” in Bach’s *C minor Partita* BWV 826).

why is the ontology of socially constructed entities different from the ontology of natural kinds? The argument from language at least tries to answer this question.

Amie Thomasson (2004; 2006) has argued that in face of the unsatisfactoriness of traditional metaphysical theories that seek to accommodate art kinds in widely recognized ontological categories (concrete particulars; abstract kinds or types; mental entities – which in turn can fall under one or other of the previous categories) a more reasonable course of action would be to revise our traditional ontological schemes to accommodate new categories that better match our pre-philosophical beliefs about art kinds (an example of this can be found in Jerrold Levinson's positing of *indicated structures* or *initiated types*, non-eternal abstract entities that are "temporally initiated" by a composer's act of "indication" in a given art-historical context) instead of forcing art kinds into ready-made ontological categories. While I'm sympathetic to this kind of openness, I doubt that metaphysical descriptivism is the only way to secure it.

The motivation for endorsing the pragmatic constraint can't *merely* be that the object of the philosophy of music are the practices of music making and appreciating, and what issues from those practices, since surely no philosopher of music would disagree with *that* (what he might contest is the idea that "the demands of the art one is attempting to understand trump the demands of metaphysics"¹³⁵ since he might disagree that the art practice of itself makes any *metaphysical* demands). Some stronger claim must motivate the pragmatic constraint. And as I see it, this stronger claim must be a claim about the *nature* of art kinds *as* socially constructed entities. This can either be the question-begging thought that we must endorse the pragmatic constraint *because* art kinds are socially constructed, or the more interesting claim based on the argument from the *qua* problem.

Now, whether or not metaphysical revisionism in the ontology of art survives the argument from the *qua* problem, it seems to me that even if it turned out that the users of art-kind terms could not be substantially wrong about the existence and identity conditions of the art kinds referred to with those terms (because, as Thomasson argues, they would have a sort of epistemic privilege about social and artifactual kinds that they lack about natural kinds), that would be compatible with meta-ontological realism (the thesis that the facts about the ontology of socially constructed things are mind-independent) precisely because facts about how cognitive agents stand to the nature of socially constructed things are *epistemic* facts. In other words, just as mountains are what they are independently of what we think of them (what depends on us is *what* things get to be counted as

¹³⁵ Kania (2008: p. 429).

mountains) socially constructed entities are what they are not in virtue of how we think of them but in virtue of how the world is (the world determines our nature as language using beings and what we are able to think and do). It would also be in virtue of how the world is that we would have, should it be true that we have it, an epistemic privilege regarding the nature of socially constructed facts and things. If Thomasson was right, we would be in a privileged epistemic position to know, e.g., what marriages are, but facts about the nature of marriages would still be epistemically objective facts, despite marriages being ontologically subjective (in Searle's sense), that is, despite socially constructed things being dependent on the thoughts and actions of cognitive agents. It doesn't follow from all this that what cognitive agents think *about* the nature of socially constructed things determines the ontological nature of those things. To take the most immediate example, no cognitive agent could make it the case that socially constructed things are not socially constructed by "thinking away" that feature of their nature. The fact that some cognitive agents think that "marriages are made in heaven" and ultimately are not a matter of contingent human convention, does not make it the case that it is so. The fact that language users believe in some version of the "natural theory of reference", expounded by Plato in his dialogue *Cratylus*, does not make it the case that word reference functions in that way. If language users have the pre-philosophical intuition that words are identical with sounds and squiggles, rather than some kind of abstract entity, their belief does not make it so, and neither do our beliefs about concepts fix the ontology of concepts, whatever these are. Our beliefs about marriages, language, words and concepts will either fit the objective nature of these things or they won't, and that is a matter of objective meta-ontological fact. To say, from the argument based on the *qua* problem, that the beliefs of cognitive agents *determine* the ontological nature of socially constructed things is somewhat like saying that a sound deductive argument for X is what makes X true. This is a misconception of what an argument is and what it does. Arguments are like telescopes that allow us to observe planets; what makes it the case that the planets have the properties they have, including being located where we observe them to be located, is not the fact that we observe them with a telescope. Arguments allow us to *see* that something must be the case, *given* other things (specified in the argument's premises). For an argument to be sound, a certain relation must obtain between the premises and the conclusion, and despite the fact that arguments are socially constructed things, whether this relation between the premises and the conclusion obtains is, once again, a matter of meta-ontological objective fact. And people are often mistaken about the nature of arguments and the relations that obtain between premises and conclusion; for instance, the confusion between the relation that obtains between the premises and the conclusion in a sound deductive argument is sometimes mistaken with whatever is the *truthmaker* of the conclusion, which are two completely distinct things. The logical and

metaphysical properties of arguments cannot be “thought away” by cognitive agents in virtue of the fact that arguments are socially constructed things. (An example of metaphysical property of arguments is their validity: for an argument to be valid it is for the world to be such that there are no circumstances where the premises of the argument are true and its conclusion false).

Suppose the world was such that no cognitive agent would be able to have false beliefs about anything whatsoever. Does it follow from here that the nature of things known was somewhat dependent on what cognitive agents think about them? Not likely. That would merely be a world where no cognitive agent would have a false belief. So, even if cognitive agents have some sort of epistemic privilege about the nature of those facts and things that are socially constructed, it doesn’t follow that the nature of socially constructed things would be determined by the beliefs of cognitive agents. The sort of privilege that Thomasson ascribes to the grounders of art-kind terms is an interesting epistemic fact, should it be true, but in no way does it endow cognitive agents with the power to determine the nature of art kinds, no more than our classificatory practices endow us with the power to determine the nature of mountains, or to think away their property of being the result of large scale geological processes. The only power we have on this matter is the power to make it the case that in our world some things get counted as symphonies, and that by “symphony” we mean a *sonata for orchestra* with certain standard, variable and contra-standard features.¹³⁶ We have the power to make it the case that some contra-standard features for symphonies become variable – composers’ decisions affect the historical *development* of musical forms – as Beethoven did when he decided to add choral parts to his ninth symphony, but neither is this a power to determine the basic ontological features of symphonies. We have no more “ontological power” over symphonies than we have over the nature of sound events.

Thomasson may be right (as I believe she is) in claiming that we won’t be able to specify the ontological nature of art kinds by applying one or other of the familiar ontological categories of *eternal abstracta*, *concrete particular*, or *mental entity*; for the simple reason that art kinds might not fit into one or other of these readymade categories. The ontology of socially constructed things may very well be a “higher order” ontology, straddling more than one of the basic categories. But this, if it is true, is compatible with meta-ontological realism: it will be in virtue of how the world is (which includes us and the way we are) that art kinds are thus, and philosophy merely tries to articulate and make those features explicit. Metaphysical descriptivism would be one way to arrive at the conclusion that art kinds straddle traditional ontological categories; another way to arrive at

¹³⁶ K. Walton (1970).

the same conclusion is by “good old” revisionist metaphysics about the nature of socially constructed things. Thomasson’s contention that cognitive agents have a sort (some degree) of epistemic privilege could be combined with a meta-ontological realist perspective. The question is whether we really have reasons to believe in such an epistemic privilege regarding the basic nature of socially constructed things (that is, whether *the world* is really such that grounders of art-kind terms cannot be substantively mistaken about the ontological nature of art kinds).

Now, as we have already seen, there are reasons to doubt that the solution to the *qua* problem requires us to endorse metaphysical descriptivism about socially constructed kinds. Dodd has argued (cogently in my view) that there are alternative hybrid theories of reference to Thomasson’s (namely, that of Gareth Evans) that do not require us to endorse a Searlean descriptivism about the “nascent ontological conceptions” associated with art-kind terms (the thesis that these “nascent ontological conceptions” disambiguate reference by *fit*). And even if we take Thomasson’s theory of reference by fit as a serious contender, I believe the previous considerations show this to be perfectly consistent with meta-ontological realism. Meta-ontological realism’s cogency is not really at stake; what is at stake, rather, is whether there are good reasons to believe grounders of art-kind terms have some epistemic privilege about the nature of art kinds.

Thomasson¹³⁷ has another argument for metaphysical descriptivism about art kinds. This argument infers the epistemic privilege of cognitive agents over the nature of artifactual kinds from the fact that the *intentions* of the makers of artifacts are partly constitutive of the artifact’s identity conditions (what this means is that we *individuate* artifacts partly in terms of their creators’ intentions). This argument, if sound, grounds a conclusion that is neither harmful to meta-ontological realism nor to the “discovery model” that goes with it. That conclusion is that

“the role of makers in establishing the nature of the artifactual kind they create endows them with some protection from certain kinds of ignorance and error about that nature, providing them with a much closer epistemic relationship to their artifactual kinds than anyone has to natural kinds.”¹³⁸

As we have seen, that we have a closer epistemic relationship with artifactual kinds than we have with natural kinds is 1) compatible with *some* kinds of ignorance and error about their nature (and *how much* error there will be is not fixed by our pre-philosophical beliefs) and 2) with the fact that our pre-philosophical beliefs do not *determine* the ontological nature of artifactual kinds. So

¹³⁷ “Artifacts and Human Concepts”, in Margolis, E. & Laurence, S. (2007).

¹³⁸ Thomasson (2007: p. 53).

Thomasson's metaphysical argument, like the argument from the *qua* problem, is not sufficient to establish her general metaphysical conclusion.

This is true for both causal-agentive artifacts and status-functional artifacts. The metaphysical feature on which Thomasson bases her argument is that of the "self-referentiality" of artifactual kinds, a feature that was also pointed out by Searle in his work on social ontology. But what is the self-referentiality of artifactual kinds? This is the feature in virtue of which, when we try to specify an artifact's identity conditions (what makes it the case that this piece of paper is money or that uttering such and such a sentence in such and such a context counts as getting married; what makes it the case that this object is a screwdriver and that object an oboe, and so on) we will have to make reference to some mental state (an intention or a belief) whose content includes a concept of the artifactual kind in question or of the function it is supposed to perform. In the case of causal-agentive artifacts, like the screwdriver or the oboe, this mental state is an *intention* to use a certain physical structure to perform some physical action (to drive screws; to produce a specific kind of sounds), an intention whose adequacy depends on an *insight* that certain physical properties will sustain certain functional properties. In the case of status-functional artifacts (like words, marriages and money), the relevant mental state is a *shared belief* or an interlocking web of shared beliefs, a collective convergence in how the object is represented, to the effect that such and such an object, in such and such a context, counts as an instance of the relevant kind. Consider the case of money: "to be represented as money" is part of the definition of money, since nothing can be money in virtue of its physical properties *alone*. This seems like a vicious regress, but actually it isn't:

"The word "money" marks one node in a whole network of practices, the practices of owning, buying, selling, earning, paying for services, paying off debts, etc. As long as the object is regarded as having that role in the practices, we do not actually need the word "money" in the definition of money, so there is no circularity or infinite regress. The word "money" functions as a placeholder for the linguistic articulation of all these practices. To believe something is money, one does not actually need the word "money". It is sufficient that one believes that the entities in question are media of exchange, repositories of value, payment for debts, salaries for services rendered, etc. And what goes for money goes for other institutional notions such as marriage, property, and speech acts such as promising, stating, ordering, etc."¹³⁹

¹³⁹ Searle (1995: p. 52).

Recall the simpler case of a territorial boundary. Imagine an early human settlement in which a causal-agentive artifact, a stone wall, is built in order to keep some people from leaving and others from entering the circumscribed space. With time, the wall decays to the point that the only thing left is a line of stones. Yet, that line of stones is still treated by the members of the settlement, as well as by outsiders, as a boundary, the crossing of which carries certain implications. The stone wall is no longer the causal-agentive artifact that it once was, since it has no *physical* power to keep people from crossing it, and neither do people have to think about it in terms of the original intention and insight that made it the causal-agentive artifact it once was. The line of stones in this example is analogous to a painting which was so severely damaged that it lost practically all its original aesthetic and representational properties. What people experience is a different physical configuration of elements that traces back to an artifact, but that is no longer the artifact it once was. Yet there is still another kind of artifact that the line of stones *is*, which is not a causal-agentive artifact but a status-functional artifact: a territorial boundary. The same physical object can coincide with both artifacts, but these are distinct. There are conceivable situations in which the causal-agentive artifact exists but not the status-functional one, and situations – as is the case with the line of stones – where the causal-agentive artifact no longer exists but the status-functional artifact still exists. This is an interesting case, since it connects to what Arthur Danto says about indiscernible duplicates and has important consequences for the theory of art: physically and perceptually indiscernible “vehicles” can be shared both by distinct works of art (think of cases, imaginary or actual, like *Pierre Menard’s Quixote*) or by an artwork and a non-artwork (think of cases, imaginary or actual, like Andy Warhol’s *Brillo Boxes*). Whether we define *art* in terms of causal-agentive artifactuality or in terms of status-functional artifactuality, one and the same item can instantiate *both* kinds of artifactuality, which means that functionalist and proceduralist theories of art *do not exclude each other*. The crux of the matter is whether one kind of artifactuality is metaphysically and explanatorily *prior* to the other, even if they can diverge to the point where we have one kind without the other. This is a point easy to miss if we confuse territorial boundaries with (causal-agentive) stone walls, a confusion which one is liable to make if one takes descriptivism about socially constructed things too seriously, should I add! Should it happen that our pre-philosophical beliefs about artifactual kinds make no distinction between causal-agentive and status-functional artifacts, descriptivism would have us incur in a false belief (assuming, of course, the Searlean story about the stone wall and the line of stones is a cogent one), and miss an important point that connects the ontology of social and artifactual kinds with a classic discussion about the nature of art.

Searle explains how the self-referentiality of status-functional kinds does not constitute a viciously regressive feature of the definition of such kinds. The concept that figures in the *definiens* is a mere place-holder for the relevant distinctions that must figure in the shared representations of the objects that fulfill the status-function in question. In the case of causal-agentive artifacts, an explicit concept for the artifactual kind in question is also not required to figure in the intentions of the artifact's maker. Plausibly, someone has made a stone axe if he has shaped a piece of stone to serve the appropriate cutting function, even if that person doesn't possess the explicit or linguistic concept *axe*. What is required is that he has certain insights (about what functional properties will be sustained by what physical properties) guiding his productive action. He will of course be able to distinguish between the different tools he makes (say, axe blades from knife blades), and so will distinguish between different objects of the same kind, but none of these requires him to have explicit linguistic concepts for the tools he makes, nor do his insights have to be articulated in a sophisticated linguistic manner, in terms of dependence relations between functional and physical properties. He merely has to see that making a wedge-wise stone shape will enable him to perform certain actions that cannot be performed with differently shaped stones. This casts a doubt on one of Thomasson's conclusions about the way intentions determine or specify the nature of artifactual kinds:

"It does seem apt to say that something is a painting only if it is the product of an intention to make a painting (as opposed to the mere unintended by-product on a drop cloth of a process of painting walls), and that something is a work of music only if it is the product of an intention to make a work of music (thus distinguishing the normal sound of people adjusting radios from that comprising a performance of Cage's 1956 "Radio Music").¹⁴⁰

Of course, whether it is actually apt to say that something is a painting only if it issues from an intention to make a painting depends on the epistemological demands one puts on such an intention. If this is the requirement that the maker of a painting must have an explicit linguistic concept for *painting* it is too restrictive, since we count cave paintings as paintings (whether or not we count them as *art*) and it is unlikely that the painters at the Lascaux or Chauvet caves had any linguistically articulated concept for *painting*. Danto (1964) argued that cave paintings are not art because their makers had no concept of art, but I don't think he would be willing to deny that they made paintings. The claim is made more plausible if by "intention to make a painting" one is not

¹⁴⁰ Thomasson (2005: p. 224).

thinking of intentions that include linguistically articulated concepts but something more close to our axe maker's ability to distinguish his axe blades from his knife blades. To make a painting, the maker of the painting must have an intention to do something *by way of* smearing pigments on a surface that is distinguishable from, say, the action of smearing the stone surface with no further purpose at all. For instance, someone produces a painting by smearing pigments on a surface with the intention to make a depiction of an animal or any other object, or with the intention of making an agreeable visual pattern, whether or not it is a depiction. The intention to make a depiction is an interesting case in point, since it doesn't require an explicit linguistic concept of *depiction*, but only the maker's (and the viewers') ability of *seeing* the depicted subject *in* the marked surface. There are such things as Paleolithic *depictions* in the absence of Paleolithic theories of depiction (which makes one wonder why shouldn't there be *aesthetically* pleasing (or repulsive) depictions in the absence of aesthetic theories or, for that matter, Paleolithic *art* in the absence of Paleolithic theories of art).

This point about the self-referentiality of artifactual kinds (both causal-agentive and status-functional) vindicates George Dickie's contention that the circularity of the institutional definition of art is not a vicious but an informative one. He argues that art-kind concepts are *inflected* so that it is inevitable that in a definition of *art*, the concepts used to define it will "feedback" into each other, as is the case with the concepts for "artwork", "artist", "artworld", "artworld system" and "artworld audience". Our adventure into the nature of socially constructed things allowed us to see that Dickie's defense of his definition is not *ad hoc*. Social and artifactual kinds really do possess the property of self-referentiality, as was explained above. But at the same time that this vindicates Dickie's strategy, it also exposes an underlying weakness: the proponent of an aesthetic theory of art can accept all of the preceding; in fact, he can accept that there is a perfectly sound *status-functional concept of art* (the concept deployed by the institutional theory) and that Dickie's definition is a perfectly apt definition of *that* concept.¹⁴¹ He can also accept that the circularity in Dickie's definition is not vicious and that it is inevitable, from what was explained above. Yet, such aesthetic theorist will also argue that the term "art" is ambiguous between two different kinds of artifactuality that objects can instantiate via their connections with cognitive agents' beliefs and intentions: causal-agentive artifactuality and status-functional artifactuality. The aesthetic theorist

¹⁴¹ See Zangwill (1994: pp. 76-78) for an argument against institutionalism that relies not on counterexamples to the institutional definition, nor on supposed epistemological weaknesses of that definition, but on the idea that "art" actually articulates two distinct concepts (and is therefore ambiguous between those concepts), one causal-agentive, the other status-functional (though he doesn't put the question in terms of "causal-agentive" or "status-functional"), which are wholly compatible, though one (that Zangwill calls "the creative concept of art") is metaphysical and explanatorily prior to the other (the "institutional concept").

can argue that Danto's cases of indiscernible duplicates are quite real, not merely philosophical fantasies, and that they must be adequately accounted for. The proceduralist approach is one attempt at making sense of them, but, so the aesthetic theorist can argue, an aesthetic theory of art has a better explanatory story to tell about them. Such cases, rather than proving the proceduralist's point, can be regimented to the aesthetic theorist's explanatory account of the nature of artworks. While it's true that indiscernible duplicates point clearly in the way of what is called in the literature a "relational account" (as opposed to theories that seek to isolate an intrinsic property or set of intrinsic properties as definitional), the contrast between intrinsic and relational is a confused contrast. Intrinsic properties are not opposed to relational properties. In fact, it is very hard to imagine what a genuinely non-relational feature of something would be; since any properties an object has it will have them in virtue of other things being the case. A less confused contrast is that of intrinsic versus extrinsic properties, and thus a truly useful distinction will be that between intrinsic and extrinsic relational properties. Accordingly, the convergences in representation on which status-functional artifactuality depends will be intrinsic to the nature of status-functional artifacts, despite such properties being relational, and the same will be true for those properties on which causal-agentive artifactuality depends, which are also intrinsic to the nature of causal-agentive artifacts. Thus, for the same (physical or abstract) object, some of its relational properties will be intrinsic when the object is considered under one sortal concept and extrinsic when it is considered under a different sortal concept. Consider the case of the stone wall, which instantiates both kinds of artifactuality – causal-agentive and status-functional – here the same object instantiates two different artifacts, whose identity conditions are different: the territorial border status-function can survive the decay of the wall into a line of stones, the causal-agentive artifactuality of the wall cannot. The same object can be considered under two different sortal concepts and in each case different relational properties will be intrinsic to its status.

A better understanding of the differences between theories of art traditionally named "functionalist" and theories of art traditionally named "proceduralist" is therefore achieved in terms of the distinction between different kinds of functional properties, already explained in section 1.1, rather than in terms of the confused contrast between intrinsic and relational.

To adjudicate between aesthetic and proceduralist theories of art, we need an account of institutional reality, and the general account of institutional reality in terms of status-functional artifactuality seems fit to the purpose. The pragmatic constraint on the philosophy of art also constitutes an attempt to provide us with a way to adjudicate between theories of social and artifactual kinds, of which theories of art (both definitional and ontological theories) are an instance. But unless the pragmatic constraint can be supplemented with an account of the mechanism by

which the pre-philosophical beliefs of cognitive agents determine the ontological nature of art kinds, it won't help us in this task, since it leaves us with at least as much difficulties as when we started our enquiry into art kinds. Namely, we have to settle on a particular version of the pragmatic constraint – what are the relevant pre-philosophical beliefs that constitute *desiderata* for a theory of art to satisfy? Different theories of art will satisfy different *desiderata* and we need a way to settle the score, a matter on which the pragmatic constraint itself won't be able to help us. If it turns out that our “nascent ontological conceptions” don't refer to anything we must settle for the “closest ontological competitor” and this poses the problem of how far can the closest competitor be for it to be a viable candidate and of how we should proceed in case there are several equally good competitors. And if we were to endorse fictionalism about artwork or musical work talk, this would leave us with the surprising result that we end up endorsing a revisionist theory, to the extent that common sense discourse about artworks and musical works is intuitively realist. As Kania (2008) put it, there is not a theory that is wholly descriptivist or wholly revisionist, while all theories on offer attempt what is called a “reflective equilibrium” between our pre-theoretical beliefs and philosophical revisionism. Now, given that the most promising argument for the pragmatic constraint – the argument from the *qua* problem – turns out not to impose a single hybrid theory of reference that gives epistemic privilege to the “nascent ontological conceptions” associated with art-kind terms, as Dodd has argued, and given that the metaphysical argument that Thomasson presents for the view that the intentions of artifact makers determine the ontological nature of artifacts is insufficient to establish her conclusion (the meta-ontological realist can accept that cognitive agents have a closer epistemological relation to the nature of artifactual kinds than to the nature of natural kinds), we cannot expect much more from the pragmatic constraint than a methodological qualm that no theorist denies anyway – which is the idea that our interest is in the actual practices of art making and appreciation. There is no super-pragmatic constraint to which we can appeal in order to solve these disagreements (e.g. who among us is as most faithful as possible to the pragmatic constraint); such disagreements must be solved through good old metaphysical theorizing.

A further point against the pragmatic constraint was summed up by Dodd in his recent discussion of Thomasson's argument from the *qua* problem against meta-ontological realism, and has to do with what the practitioners in the relevant art practices expect from philosophical theories of art: “Largely uninterested in explaining how works of music are repeatable and audible in performance,

practitioners simply adopt, quite unreflectively, an ontological conception of such works that they assume to be a corollary of their appreciative practice.”¹⁴²

While nothing prevents artists from being apt philosophical theorists of their own practice, it is generally true that they are above all concerned with *making* art, not with solving philosophical puzzles raised by their practices. Take, for instance, Joseph Beuys’ definition of art: “Art is life, life is art.”¹⁴³ This definition is very easily defeated: we either take it literally, in which case it is obviously false, since it has the consequence that my cat meowing is a work of art, for it is a living thing, while the Victory of Samothrace is not a work of art, for it is not a living thing; or we interpret it as a figurative and rhetorical expression of a point which is only faintly related to the definition of art, namely, that art and aesthetic concerns are pervasive in human life, not just on a few established art forms, and that we can find artistic elements in many activities that aren’t standardly classified as art.

Another case in point are Frank Zappa’s insightful remarks on the function of The Frame in art, which amount to an intuitive formulation of a proceduralist theory of art (whether Zappa was writing with knowledge of Danto and Dickie’s theories, I have no clue).

The most important thing in art is The Frame. For painting: literally; for other arts: figuratively — because, without this humble appliance, you can’t know where The Art stops and The Real World begins. You have to put a ‘box’ around it because otherwise, what is that shit on the wall? If John Cage, for instance, says, “I’m putting a contact microphone on my throat, and I’m going to drink carrot juice, and that’s my composition,” then his gurgling qualifies as his composition because he put a frame around it and said so. “Take it or leave it, I now will this to be music.” After that it’s a matter of taste. Without the frame-as-announced, it’s a guy swallowing carrot juice. So, if music is the best, what is music? Anything can be music, but it doesn’t become music until someone wills it to be music, and the audience listening to it decides to perceive it as music.”¹⁴⁴

The “frame” here isn’t essentially a material object but a *procedure* that may or may not be signaled through a material object. Zappa’s account is sophisticated enough not to make it all depend on the intentions of the artist; it puts an extra constraint on the artworld audience’s acceptance of the would-be artwork as a candidate for appreciation. It is only by being accepted into an appropriate institutional frame by way of appropriate procedures that something becomes a work

¹⁴² Dodd (2012: p. 79).

¹⁴³ Stachelhaus (1991: p. 1)

¹⁴⁴ Zappa & Occhiogrosso (1990: p. 140).

of art (John Cage can make it the case that his gurgling of carrot juice counts as a composition or a “candidate for appreciation” because as an artist he has the requisite “authority” to bring the conventions of the artworld to bear on his composition). Yet one can reasonably ask if this is a good account of *all* art, even if it is a good account of *some* art. Is it the case for all paintings that they are not art unless they are “framed” as art? Such theories have a hard time accounting for the first artworks, which could not have an institutional frame into which they could be accepted. The proceduralist is forced to see the art status of first art as something that those objects acquire retrospectively. Proceduralism also has difficulties accounting for art produced outside the boundaries of the artworld, which seems not only a very real possibility but something quite pervasive. There is more art in the world than can be accounted for with the concept of artworld art. Now, though I believe Zappa would surely have interesting things to say about these questions, we would not evaluate them with an appeal to the pragmatic constraint, since other artists would certainly have other interesting things to say while disagreeing with Zappa. And I doubt that any artist would be able to present interesting answers to these questions making use of no more than their pre-theoretical beliefs and intuitions. These are simply not enough and as soon as one begins to think about the subject, one is no longer relying merely on intuitions. Besides, there is nothing magical about intuitions; intuitions, like the rest of our beliefs, are defeasible.

If the pre-theoretical beliefs and intuitions of artists are of no avail in matters of definition, why should they constraint philosophical theorizing when it comes to the ontological nature of art kinds? The problems we have in adjudicating between different philosophical views of art kinds will be paralleled by the problem of adjudicating between conflicting beliefs of artists about art kinds. In particular, to what extent “pre-theoretical” beliefs are really “pre-theoretical”? And if the beliefs of artists will inevitably be “theory laden”, how much “contamination” will be acceptable before we stop considering a given belief as relevant data for the philosopher of art? Once again, there is no super-pragmatic constraint to answer these questions. The discussion is philosophical through and through.

This is not an argument for the exclusion of artists’ beliefs from philosophical consideration, of course – quite the contrary. What happens is that all relevant beliefs will have to be evaluated philosophically, and consistency with common sense, with pre-theoretical practice, etc. is but one epistemological virtue to be considered in the evaluation of a philosophical theory. It is not the sole or the most important epistemological virtue. Most importantly, it is not a way of extracting a ready-made theory of art or an ontology of art kinds. Unless a better argument is presented for metaphysical descriptivism in the philosophy of art, some cogent description of a mechanism by which the pre-philosophical beliefs of artists determine the ontological nature of art kinds, we are

stuck with the good old method of evaluating philosophical theories: by argument, counterexample and thought-experiments.

While any philosopher of art will agree that our interest in doing the philosophy of art is to try to understand *the practices* of art making and appreciation, and not what any philosopher would *like* those practices to be in order to fit better with his general philosophical outlook, does a serious commitment to the pragmatic constraint compel us to take one or other theoretical road? I mean, do the *actual* practices of artists and critics make any serious metaphysical *demands* of us, over and above a “nascent ontological conception” so general that the proponent of any metaphysical theory can argue his way into a plausible consistency with such conception? In other words, do the actual practices of artists and critics constrain the philosopher’s theorization beyond what can be already summarized in a list of those puzzling features of artworks that invite philosophical reflection in the first place? Such a list would be, more or less, something like the following: a) that (at least some) musical works are *performable*; b) that they can have many instances; c) that performances can have a more “thick” array of properties than what is specified in instructions for performance; d) that there is a difference between versions of a work and distinct works that are based on the original work. Granted, we also have more troubling pre-theoretical intuitions: that artworks are created by their authors (which seems to clash with the “intuition” that many works are multiply instantiable); that multiply instantiable works are as temporally and modally flexible as singular works (which seems to clash with the “intuition” that multiply instantiable works must be some sort of abstracta); that paintings and non-cast sculptures and similar works are singular physical artifacts; and so on. But we have troubling intuitions about any proper subject of philosophical reflection. The point of doing the philosophy of a given domain is to attempt to clarify and resolve the puzzles that arise from those intuitions. If our actual practices of art making and art appreciation really do pose any metaphysical demands on the philosopher of art, those demands will in turn have to be philosophically evaluated.

3 A working notion of the aesthetic

That we can lose things, that in fact we are always in the process of losing everything we have, underlies the longing with which we inhabit the world. And in that longing resides the possibility of beauty.

Crispin Sartwell¹⁴⁵

What makes a theory an “aesthetic” theory is its deployment of aesthetic concepts. But what is meant by an “aesthetic concept”? First, the most general of such concepts is the very concept of the aesthetic itself, that is, the concept of a particular domain of reality; second, aesthetic concepts are the concepts of all that we discriminate within that domain, that is, all those things that fall under the concept of the aesthetic: aesthetic judgment, aesthetic experience, aesthetic predicate, aesthetic property, aesthetic value, aesthetic function, aesthetic intention, aesthetic discrimination (taste), aesthetic pleasure, and so on. Acknowledging these distinctions doesn’t compromise us with the idea that *any* aesthetic concept (including the concept of the aesthetic itself) is a truly useful concept, that it is “in good shape” or that it picks out what it purports to. For instance, someone may accept that there are aesthetic properties while being skeptic about aesthetic experience, i.e., she may think that the concept of aesthetic experience fails to pick out a distinctive *kind* of experience, even though she believes we do experience aesthetic qualities of objects and that such experiences ground our aesthetic judgments. Likewise, someone may be skeptic about the notion of aesthetic pleasure, that is, a pleasure which we would distinctively feel or introspectively identify as “aesthetic”, as opposed to other kinds of pleasure, without being skeptic of aesthetic properties or of the normativity of aesthetic judgments. One can also accept the idea that aesthetic judgments are normative while not accepting that things have aesthetic properties or qualities (two examples are Kant and Hume). Furthermore, someone who endorses the idea that things have aesthetic properties has enough latitude to choose between competing ontological theories of such properties. Between the extremes of aesthetic realism (the idea that there are mind-independent aesthetic facts) and uncompromising aesthetic eliminativism there are other options on offer.

I don’t intend here to defend a particular view of aesthetic properties, regarding their ultimate ontological status (i.e. I don’t wish nor do I need to make a choice between realism and anti-realism about aesthetic properties), though I must give some characterization of what I take the experience of aesthetic properties to be, of what it is to experience reality as endowed with aesthetic properties,

¹⁴⁵ (2004: p. 4)

since the concept of aesthetic property will figure ineliminably in my explanation of what music is. Yet, this fact does not commit me to either a realist or anti-realist theory of aesthetic properties.

As it happens, the view I consider most plausible lies somewhere in between the extremes of realism and anti-realism (i.e. aesthetic properties are located neither in the outside world *tout court*, nor are they merely “in the eye of the beholder”); a view according to which aesthetic properties are partly mind-dependent, and partly dependent on objective properties of objects, including some of their contextual properties (what we may call the *dependence base* of aesthetic properties). This is a view that makes room for aesthetic normativity, but at the same time is not committed to the idea that all aesthetic judgments must be equally objective. A cognitivist about aesthetics is not thereby committed to the view that all aesthetic evaluations are objective, but rather to the view that not all of them are subjective.¹⁴⁶ Having said that, I don’t want to set as a condition for my account of music the acceptance of any particular doctrine concerning the precise nature of aesthetic properties. There should be enough leeway for realists and anti-realists alike to adapt an aesthetic theory of music (or of art, though one does not need to endorse such a theory about both) to their own favored account of what aesthetic properties are and to different conceptions of where the source of aesthetic normativity lies.

What is it to experience reality as endowed with aesthetic properties? Though there isn’t a universal consensus as to what properties count as aesthetic (or, to put matters another way, as to what predicates genuinely perform an aesthetic descriptive role, whether or not they describe real aesthetic properties), there has been enough agreement to make the notion of an aesthetic property intelligible. The two central aesthetic properties are *beauty* and *ugliness*. In the eighteenth century, philosophers added the concept of the *sublime* to the list. Though nowadays it is quite common to use “sublime” as synonymous with “superlatively beautiful”, its original meaning was connected to the awesome or awe-inspiring, the wondrous, the immense, that which challenges or exceeds human imagination or understanding, that suggests or hints at the infinite or the unfathomable, in sum: that which dwarfs us. Accordingly, things such as a mountain range, massive waterfalls, the starry sky, the sight of a natural catastrophe, are appropriately called “sublime”. The qualities referred to with the predicate “sublime” are thus significantly different from the qualities we refer to when applying

¹⁴⁶ There is a potential ambiguity in “subjective” here that needs clarification. I just used the expression “subjective” as applied to evaluations or judgments in the sense of lacking normativity or standards, i.e., of being wholly subject-relative. In Kantian terminology, a “subjective judgment” is a judgment made on the basis of a felt response. In this sense, both judgments of beauty and ugliness and judgments of niceness and nastiness are “subjective”, but Kant distinguishes the former as having a normative base while the latter don’t. So, the ambiguity is between a judgment whose truth or aptness is *relative to the subject’s response*, and a judgment that is made on the basis of a felt response to an object (as opposed to a judgment made on the basis of testimony).

the predicate “beautiful” to things such as a rose, a bonsai tree, the nightingale’s “song”, human faces, or a polished chunk of lapis lazuli. However, there is an ambiguity in the notion of beauty allowing us either to count sublimity as *one of the ways* in which something is said to be beautiful, or as a quality quite distinct from a beauty-making quality. This ambiguity can be clarified via a distinction which is of key importance to contemporary aesthetics, namely, that between *general* or *verdictive* aesthetic properties and *substantive* aesthetic properties (Zangwill: 2002). The same distinction can be phrased in terms of “thin” and “thick” terms: thin aesthetic terms are those that have little or no “descriptive content”; these are the most general terms of aesthetic appraisal, such as *beautiful* and *ugly*, but also *unified*, *coherent*, *balanced*, *harmonious*. Thick aesthetic terms, on the other hand, are said to be richer in descriptive content: *elegant*, *graceful*, *delicate*, *dainty*, *dumpy*. (The idea that such terms are more “descriptive” comes from the fact that, e.g., “elegant” excludes more from its extension than “beautiful”, since all elegant things will be beautiful to the extent of their elegance, while many beautiful things are not elegant. In other words, there are more ways of being beautiful than of being elegant, which is merely *one* of the ways for something to have beauty.) Concerning beauty, the ambiguity lies in the fact that “beautiful” is sometimes understood as a thin predicate, in which case it is synonymous with “having positive aesthetic value” (correspondingly, the predicate “ugly” is synonymous with “having negative aesthetic value”), and sometimes it is understood as a thick predicate (denoting a kind of aesthetic quality ascribed to roses, human faces, bonsai trees, or polished chunks of lapis lazuli, but *not*, e.g., to things appropriately termed “sublime”). Used as a thick predicate, “beautiful” is somewhat more close to what we mean by “pretty”, even if this isn’t the appropriate word for every case (“charming”, “gorgeous” and “delightful” are among the likely alternatives). Incidentally, there seems not to be an equivalent ambiguity for “ugly”, so that by using it we invariably mean to ascribe negative aesthetic value in general.

The problem of using “beautiful” as a predicate for positive aesthetic value in general is that there are cases where we want to say that something is aesthetically valuable and yet we wouldn’t describe that thing as beautiful (examples of “sublime” things, in the original sense of the word, are precisely a case in point). For instance, certain figurative paintings may be aesthetically impressive for reasons other than being beautiful or visually pleasant; in fact, they may be depictions of things we would call “ugly” and yet the images are not ugly in the sense of being aesthetically worthless, unappealing or ineptly made. However, bearing this ambiguity in mind plus the context of utterance is enough to understand when the word “beautiful” is being used in the verdictive sense and when it is used in the substantive one.

In the twentieth century, the list of recognized predicates purportedly denoting aesthetic properties has widened immensely.¹⁴⁷ The *locus classicus* for the contemporary discussion of aesthetic properties is Frank Sibley's 1959 paper, "Aesthetic Concepts":

"Aesthetic terms span a great range of types and could be grouped into various kinds and subspecies. (...) Their almost endless variety is adequately displayed in the following list: *unified, balanced, integrated, lifeless, serene, sombre, dynamic, powerful, vivid, delicate, moving, trite, sentimental, tragic*. The list, of course, is not limited to adjectives; expressions in artistic contexts like "telling contrast", "sets up a tension", "conveys a sense of" or "holds it together" are equally good illustrations."¹⁴⁸

Other aesthetic predicates (or predicates that can be put to an aesthetic use) mentioned in Sibley's paper are: *graceful, dainty, handsome, comely, elegant, garish, melancholy, lovely, pretty, beautiful, flamboyant, majestic, grand, splendid, fiery, massive, flaccid, weakly, washed out, lanky, anaemic, wan, insipid, joyous, robust, strident, turbulent, gaudy, chaotic, strong, powerful, monotonous, forceful, gay, taut, solemn, ugly, exquisite, nice*.

The "almost endless variety" of such terms suggests there is no definite or closed set of specifically aesthetic predicates, since virtually any term or expression can be conscripted for aesthetic use. Conversely, aesthetic terms can be used non-aesthetically, as in "a beautiful hand at cards". A more controversial case is that of what Collin McGinn refers to as "terms of moral appraisal that have a strongly aesthetic flavor"¹⁴⁹: *fine, pure, stainless, sweet, wonderful, rotten, vile, foul, ugly, sick, repulsive, tarnished*. The reason I say this is a controversial case is that it is an object of dispute whether descriptions such as "an ugly action" and "a beautiful character" are cases of metaphorical, non-aesthetic use of aesthetic terms, or rather cases of genuine ascription of aesthetic properties. Can a predicate be used morally and aesthetically at the same time? Are some moral evaluations aesthetic and vice-versa? I call attention to such cases in order to point out that in all likelihood there is not a clear-cut boundary between the aesthetic and the non-aesthetic use of predicates, which of course doesn't mean that there aren't clear central cases of each.

Sibley's own account of what is common to all predicates used aesthetically is that their application, as opposed to predicates like "is square", "is blue" or "is five feet tall", requires the

¹⁴⁷ Rafael De Clercq, in his paper "The Structure of Aesthetic Properties" (2008: p.895), provides a useful table listing the predicates considered by several authors to denote aesthetic properties (or, to phrase it in a way more congenial to anti-realism: predicates considered to function aesthetically).

¹⁴⁸ Sibley (1959: pp. 421-422).

¹⁴⁹ *Ethics Evil and Fiction* (1992: p. 92)

exercise of *taste* or aesthetic discrimination. This view has been attacked on the grounds of its circularity, since taste, in this context, will in turn be defined as that which allows us to recognize aesthetic qualities.

It is difficult to pin down what all terms, insofar as they are used aesthetically, have in common. I speak here of an aesthetic *use* of terms or predicates, rather than simply *aesthetic terms*, since few of the predicates used to describe the supposed aesthetic qualities of objects have their primary application (or their literal meaning) in the aesthetic domain. “Beautiful”, “ugly”, “elegant”, “garish” and “graceful” are examples of predicates whose primary use is aesthetic; but terms such as “delicate” or “serene” are not used aesthetically when applied literally to an object. A vase can be delicate in the quite literal sense that it is liable to break into pieces if dropped; but when applying the term aesthetically we are not referring to the object’s fragility or its disposition to break if dropped. Applied aesthetically, predicates such as “delicate” or “serene” are used non-literally, to describe certain features of how an object *looks* or *sounds*. When we say of a painting that it is “balanced” or “vibrant” we do not mean literally that it is in equilibrium with some other object or that it quivers. If we describe a melody as “delicate” we do not mean that it can break if dropped, since melodies are made of sounds and sounds can’t break nor be dropped, for that matter. Or, to pick a specific musical example, if someone describes Beethoven’s piano sonata in B Flat Op. 106 (*Hammerklavier*) as “craggy”,¹⁵⁰ she does not mean that the music literally has many crags.

A temptation here is to say that aesthetic terms or concepts are value-laden or that they have built-in evaluative direction. But this is not always the case. Think of descriptions of music in emotion terms, such as “sad”, “happy”, “somber” or “mournful”. That such descriptions are appropriate tells us nothing as to whether the music is good or bad. Nevertheless, it is true that, all things being equal, we take qualities such as balance to be good-making qualities. Yet, consider the case where something is so unified and coherent to the point of being boring or dull, or so diverse to the point of being cacophonous, noisy or unintelligible. A different kind of example is when a given aesthetic quality in an object conflicts with the object’s function or purpose so as to seem out of place, awkward or making the design incoherent.¹⁵¹ Hence, there are cases where otherwise good-making qualities interact in such a way that the result is anything but good. In view of these counterexamples, besides the fact that many value-laden terms, such as “great”, “lousy”, “terrific”

¹⁵⁰ The example is from Levinson (1980: p. 17).

¹⁵¹ Nick Zangwill (2001: p. 18) gives the example of delicacy in a triumphal arch, which is perhaps out of place.

or “marvelous”, typically do not function aesthetically,¹⁵² evaluative directedness is arguably not the best way to distinguish aesthetic terms and the qualities they are supposed to refer.

One way of characterizing what it is to experience reality as having aesthetic properties that I find heuristically useful was presented by Eddy Zemach in his book, *Real Beauty* (1997).¹⁵³ Though Zemach espouses a particularly strong version of aesthetic realism, the feature of his theory that I highlight next is acceptable for realists and anti-realists alike, since it is a way of conceiving of what it is to experience reality as endowed with aesthetic properties, without committing one to any specific ontological doctrine about the reality and nature of aesthetic properties.

In rough outline, the idea is the following: cognitive agents as ourselves are organisms capable of two important kinds of mental state, namely, *beliefs* and *desires*. As beings capable of forming beliefs and endowed with a perceptual apparatus, we are able to discern phenomenal properties of things, such as their being blue, square, etc. As beings capable of desires, that is, as not only capable of representing the world but also being interested in how the world affects us, we also inevitably *interpret* the world as impacting our sensibility.

“A mind that has no interests can discern non-aesthetic phenomenal properties, say, see X as blue, but it cannot see things aesthetically, that is, as having aesthetic properties. Aesthetic properties appear only to those whose seeing is modulated by desire.”¹⁵⁴

Consider a few simple examples. We see things as colored, and we also perceive colors as *warm* and *cool*. These terms are not applied literally, since no red coated surface needs to be actually warm for the color to be perceived as warm, nor does a blue coated surface need to be actually cool for the color to be perceived as cool. The predicates “warm” and “cool” are being used to qualify visual sensations. Some of the most basic cases of aesthetic description are perhaps that of terms used to describe cross-modal sensory perception, as when predicates are applied to the way a thing looks or sounds that have their primary application to other sense modalities. The reason I count these cases among the “most basic” ones of aesthetic description is that they are probably less subject to cultural conditioning (not that degrees of cultural conditioning are incompatible with a description qualifying as aesthetic; in fact, a degree of cultural conditioning may be required for one to apprehend certain aesthetic features of things, without the resulting experience being any less

¹⁵² See Kivy (1975: p. 197).

¹⁵³ See, in particular, chapter 5, “The Ontology of Aesthetic Properties”.

¹⁵⁴ Zemach (1997: p. 103).

perceptual on account of it)¹⁵⁵. Perhaps even more basic examples are found within the same sense modality. Consider the difference between the description of a surface as “smooth” or “coarse” as opposed to its description in more purely physical terms, such as of offering more or less resistance to movement. Here we draw closer to the boundary between aesthetic and non-aesthetic description, in the sense that terms such as “smooth” and “coarse” are seemingly more descriptive of the non-aesthetic properties in virtue of which the aesthetic use of those terms might be appropriate. A smooth surface is even, uniform, with no irregularities; but the way “smooth” applies literally to a surface differs from the way “delicate” applies literally to an object, in that the literal, non-aesthetic use of the latter merely refers to the object’s propensity to break under certain conditions; “smooth”, on the other hand, even when literally applied, carries a degree of “psychological contamination”, i.e., it carries information that is not merely descriptive of the physical properties of the object, but also of how those properties impact on our sensibility. To describe something as smooth is also to describe certain properties of it as appealing to us in a peculiar manner, as somehow responding to our expectations. The use of the predicate “smooth” is even more obviously aesthetic when applied to visual appearances of objects and to sounds. When speaking of the smoothness of a musical line, we’re not referring to tactile properties of sounds and yet it seems that acquaintance with the tactile properties that ground the ascription of smoothness to a surface is somehow necessary for one to grasp the concept of musical smoothness.

Here it might be objected that all perceptual concepts “carry information” about the way the physical properties of objects impact on our sensibility, so that my previous observation regarding smoothness is not particularly informative about what distinguishes aesthetic qualities from sensory non-aesthetic ones, such as colors, smells, sounds, etc. However, there is an important difference. Take the example of an object said to be *garish*. If asked to justify our judgment that the object is garish we would probably point out that it is overly colored, bright and glaring. Now, while we can enumerate the object’s perceptual properties – its several colors by hue, saturation, brightness, etc. – independently of each other and of the object’s other properties, we cannot conceive its garishness independently of its phenomenal properties. We could perceive a colored patch of that object in isolation from the others, but we can’t perceive its garishness apart from perceiving its other properties, particularly its color properties. The point here is that an object cannot be garish except in virtue of *other* features it has; its non-aesthetic perceptual properties (or some of them) are what

¹⁵⁵ Christopher Peacocke gives a parallel example with the perception of sounds as meaningful linguistic utterances: to perceive sounds *as words* and to perceive them as meaningful sequences, one needs to have learned a language, but the resulting experience is no less perceptual on account of it. See the podcast interview: <http://podbay.fm/show/323127331/e/1307973600?autostart=1>

makes it garish. Arguably, many other facts are required to explain the garish quality of the object, namely, facts about our perceptual and cognitive make-up; contextual facts regarding the relations between the object and its surroundings and relations to other objects; the concept of garishness (or garishness associated with gaudiness) may even include moral hues, in that it suggests lack of restraint or vulgarity, at least in cases where garishness is predicated of human artifacts. So, while all perceptual concepts carry information about how an object's properties affect us (we can't conceive the color of an object independently of our *visually experiencing* that object as a colored object), aesthetic concepts differ significantly from perceptual concepts in that the qualities they purportedly denote can't be perceived independently from the object's other properties, in the sense that a colored patch of a garish object's surface can be perceived separately from a differently colored patch of the same object. An aesthetic quality of an object is not just one more perceptual property of that object, yet it can't be detached from its perceptual properties and, arguably (even if not always), from some of its contextual properties (in order to be seen, in fact, in order to be *present*, qualities such as garishness require contrast against some background). One can say that aesthetic properties *emerge* from non-aesthetic properties (arguably, from *both* perceptual and contextual properties)¹⁵⁶, not being reducible to them, in the sense that we cannot *explain* what it is for something to have a certain aesthetic quality (or what a certain aesthetic predicate means) by listing a set of non-aesthetic qualities (or a set of non-aesthetic descriptions) that would be either necessary or sufficient for the possession of the aesthetic quality (or the applicability of the aesthetic predicate). To explain *in purely non-aesthetic terms* what it's like to experience something has having an aesthetic quality will be at least as difficult as explaining the smell of coffee to someone who has no idea what the experience is like. If two individuals experience the same properties of a painting as visually "clashing", then both will understand the expression "visual clash" without being given a description of the painting's perceptual properties and independently of their ability to describe what they experience. Frank Sibley put the matter in terms of aesthetic concepts' being merely *negatively* condition governed. What this means is that given some true non-aesthetic descriptions of an object, these may be incompatible with the application of certain aesthetic predicates (e.g. consisting of only two bars of very pale blue and grey is incompatible with

¹⁵⁶ Taking that contextual properties are a subclass of relational properties, and that relational properties can be either perceptual or non-perceptual, four positions about the dependent character of aesthetic properties present themselves: a) aesthetic *formalists* assert that aesthetic properties emerge from (or depend on) perceptual properties *alone*; b) *moderate formalists* assert that aesthetic properties *can* emerge from (or depend on) *both* perceptual and non-perceptual properties but *not* on the latter alone. The other two alternatives in logical space are c) that aesthetic properties *always* emerge from (or depend on) *both* kinds of non-aesthetic properties and d) that aesthetic properties *can* emerge from (or depend on) non-perceptual properties alone.

being garish; a counterfactual situation in which “garish” is appropriately applied to a painting consisting only of two bars of very pale blue and grey is a situation where the word “garish” is used to denote a different quality from what it actually denotes). In other words, while there are sufficient conditions *not* to apply a given aesthetic predicate, no set of non-aesthetic descriptions is either necessary or sufficient for the application of any given aesthetic predicate. This doesn’t mean that *no* descriptions can be offered as reasons for applying an aesthetic predicate; only that such reasons are neither necessary (the same predicate can apply in virtue of different non-aesthetic properties) nor sufficient (something can have all the features listed and still lack the aesthetic quality). Furthermore, the *same* aesthetic property can be exemplified by completely different things, in an endless variety of ways. Gracefulness in a painting is very different from gracefulness in music and both differ from gracefulness in nature.

Even though these have been described as such, I don’t think brightness and shininess can be appropriately described as aesthetic properties. At most, they belong in the dependence base of aesthetic properties, and the predicate “bright” may be used as an aesthetic predicate in some of its non-literal applications. “Intensely coloured” is perhaps a borderline case, since intensity is one of the “canonical” aesthetic properties. Briefly, the rationale for this is that aesthetic predicates cannot be replaced by descriptions of perceptual properties without loss of meaning. This is because someone may see that an object has saturated bright colours and fail to see that it is unbalanced, garish, gaudy and the like.

Zemach’s insight about the conceptual connection between aesthetic qualities and desire, I believe, helps us to shed light on some of the above-mentioned phenomena, more than rival accounts that seek to capture what is specifically aesthetic about an experience by focusing on supposedly distinctive hedonic, epistemic or axiological¹⁵⁷ features of that experience. And in order to explore the conceptual connection between aesthetic qualities and desire one need not buy the whole package of Zemach’s theory, namely his unyielding realism about aesthetic properties. An aspect-perception theorist about aesthetic qualities (someone who holds that aesthetic judgments are not, strictly speaking, representations of properties in objects but are rather based on imaginative experiences) could accommodate this idea perfectly well.

But how do we unpack the idea that aesthetic qualities are conceptually connected with desire? First of all, we need to bear in mind that this is not the same as identifying aesthetic qualities with

¹⁵⁷ Respectively, views according to which “aesthetic experience” is defined by 1) a distinctive pleasure quality, singled out as “aesthetic pleasure”; 2) a distinctive *mode of attention* to objects, such as “disinterested attention”; 3) the *intrinsic value* of an experience, or its being *valued for its own sake*.

feelings of desire or repulsion. They are not to be identified with any pleasurable/unpleasurable feelings that we somehow “project” onto objects. The idea here is to situate the aesthetic in the virtuous middle between two implausible extremes: 1) a view that simply identifies aesthetic qualities with hedonic responses (call this the “liberal” view); and 2) a view that identifies all aesthetic qualities with narrow, cognitively oversophisticated, culturally conditioned responses to objects (call this the “narrow” view). In his recent book, *The Artful Species*, Stephen Davies argues that these have been the two extreme accounts of the aesthetic generally assumed, respectively, by several biologists and psychologists (the liberal view) and philosophers of aesthetics (the narrow view). We should avoid both extremes. The liberal view makes the concept of the aesthetic amorphous and redundant. The narrow view makes the concept arbitrarily restrictive, it obscures the actual plausibility of aesthetic theories of art that don’t rely on the narrow view (by obscuring the fact that the narrow view itself is not a necessary condition for an aesthetic theory), as well as providing a straw man for skeptics about the aesthetic. We find the ultimate expression of the narrow view perhaps in the often quoted line from Clive Bell’s *Art*: “to appreciate a work of art we need bring with us nothing from life, no knowledge of its ideas and affairs, no familiarity with its emotions.”¹⁵⁸ The position roughly delineated in this sentence is what has come to be known in the literature as *formalism*. Its champion in the specific case of music is the nineteenth century critic Eduard Hanslick, who famously observed that music is nothing but “tonally moving forms” and held that beautiful *organization* of sounds (independently of any extra-musical ideas) is the ultimate aim of music making and the only appropriate object of music appreciation. In making his case against both the liberal and the narrow view, Davies traces the latter back to the Kantian theory of the aesthetic, in particular the notion of disinterested contemplation.

I recall here Augustine’s dilemma for the aesthetic subjectivist in his work *On True Religion*: “I especially ask whether things are beautiful because they please or whether they please because they are beautiful.”¹⁵⁹ The first horn of the dilemma corresponds to the liberal view above: beauty is identified with a positive hedonic response. To say “X is beautiful” is analyzed as “X pleases me” or “X is pleasant”. This immediately faces the following difficulty: even if it’s true that all beautiful things are always pleasing, not all things that please are appropriately called beautiful. How did we learn then to discriminate between those hedonic responses that are beauty-relevant and those that aren’t? Do some pleasant sensations just *feel beautiful* while others don’t? This is an utterly obscure notion. Yet, if one says that things please because they are beautiful, one commits himself to the

¹⁵⁸ Bell (1914: p. 27).

¹⁵⁹ Quoted in Tatarkiewicz (2006: p. 59).

view that beauty is independent of hedonic response, and the intuitive alternative is that beauty is a mind-independent property of things. Kantian aesthetics was an attempt to navigate between the two horns of the dilemma, denying that attributions of beauty are disguised ways to referring to the hedonic response itself and also denying that beauty is a mind-independent property of objects. In Kant's theory, aesthetic attributions (judgments of taste) *depend* on hedonic responses (that is what makes them *subjective*) but the hedonic response is not what *makes* the object beautiful; pleasure in beauty, in this view, is pleasure taken in the way the object appears to us in a certain kind of experience, where the cognitive faculties of the understanding and the imagination function in an unusual way – this is what Kant called the “free play” of the cognitive faculties, which would characterize aesthetic experience, distinguishing it from ordinary experience. This notion of free play is somewhat obscure and it generates some problems for the Kantian view – if beauty is not a property, why do some objects give rise to this peculiar functioning of the cognitive faculties while others don't? Kant tries to supplement this with the notion of a “form of purpose” exhibited by those objects that trigger the free play of the faculties. But it is not clear how this notion can be helpful, since anything can be said to exhibit “form of purpose” (particularly, things we would describe as “ugly” can be said to share this attribute or be experienced as sharing it, so this feature cannot be what distinguishes beautiful objects). But if no real properties of the object are partly responsible for our experiencing the object as beautiful, then the experience is wholly arbitrary. Anyway, I'm not concerned here with the interpretation or soundness of Kant's view, and I merely mention it in order to map the conceptual alternatives in thinking about the experience of aesthetic qualities.

Whatever the mechanism is that accounts for our experiencing things aesthetically, the relevant fact for the aesthetic theory of music is that we *do* experience things that way. Whether or not beauty and ugliness (and all the other qualities referred to with aesthetic predicates) are *real* properties of objects – independently of how they are to be ultimately conceived – there is a *phenomenal* difference between experiencing something as having this or that aesthetic quality; there is such a thing as *what it's like* to experience something as elegant, in contrast with what it's like experiencing something as having a quality incompatible with elegance. We distinguish between correct and incorrect uses of aesthetic predicates as referring to certain qualities, even when we disagree about a given object's having the qualities in question; that is, we disagree about whether a particular object is elegant, but we agree on the distinction between elegance and other aesthetic qualities (e.g. we agree on what elegance *excludes*, even if we disagree about the presence or absence of the excluded qualities). In fact, this is a necessary condition on *disagreeing* about whether X is elegant; if “elegant” *meant* different things for different people there could not be any

genuine disagreement on aesthetic matters, in which case disagreement could not be invoked to ground a relativist stance in aesthetics. And neither can the mere existence of disagreement by itself ground such a relativist stance, since there is disagreement about all kinds of property attributions, not only aesthetic ones; for instance, disagreement about what is true does not by itself yield relativism, and disagreement about whether an object has one or other particular shade of color does not yield color “projectivism” or subjectivism. Color perception, as is the case with other non-aesthetic perceptual properties, crucially depends on what is called “standard conditions of observation”, which includes facts about the perceivers’ perceptual apparatus.

The widespread agreement on the meaning of aesthetic predicates, people’s ability to learn how to properly use aesthetic predicates in the absence of any standard behavior for aesthetic appreciation (there is a typical behavioral pattern for experiences of pain, but what could count as a behavioral pattern for finding something elegant?) was obscured by the fact that earlier aesthetic theorists focused almost exclusively on the *verdictive* properties of beauty and ugliness. But while people may agree on substantive aesthetic descriptions of a particular object, they can disagree about the relative aesthetic value of that same object within a comparison class. Furthermore, there is also a tendency to exaggerate the extent of disagreement on aesthetic matters. The eighteenth century philosopher Edmund Burke has noted this tendency, pointing out that agreement in such matters is often more common than the popularity of the dictum *de gustibus non est disputandum* (rivaled in our own time by *beauty is in the eye of the beholder*) suggests:

“And indeed on the whole one may observe, that there is rather less difference upon matters of Taste among mankind, than upon most of those which depend upon the naked reason; and that men are far better agreed on the excellence of a description in Virgil, than on the truth or falsehood of a theory of Aristotle.”¹⁶⁰

Of course, one should also beware of exaggerating in the opposite direction, overly downplaying the extent of disagreement in aesthetic matters. But it must be acknowledged that disagreement implies a measure of agreement – there must at least be agreement on what one is disagreeing about; to deny an ascription of elegance or delicacy one must agree on what the disputed quality is (independently of what particular ontology of aesthetic properties one endorses). One must agree on *what it is like* to experience an object O as having aesthetic quality A in order for there to be disagreement about ascribing A to any particular object. Whether there is anything more to A than

¹⁶⁰ Edmund Burke (1990: p. 23).

the *experience* of O as A is a distinct question, over and above our ability to genuinely disagree on aesthetic matters.

The fact that disagreement about possession of A by O can't be a matter of disagreeing about the *meaning* of predicates for A was noted by Benjamin Tilghman, in a passage that deserves to be quoted in full:

“One person says this is fustian and someone else says no, it is not, but this other is fustian. It is tempting to suppose that in a case like this there are two conflicting opinions about the meaning of the word “fustian,” one thinks that it means a thing like this, while the other believes that the word means a thing like that. Here we may be inclined to say that there is no real dispute between the parties and to take a sometime popular “semantic” stance and dismiss the dispute as merely a quibble about words. The implication seems to be that agreement about the meaning of the words would lead to agreement about the truth of the judgments. That, however, cannot be the way to look at it. It is surely not right to say that the critics had affixed different *meanings* to their expressions; it is likely all would agree with the dictionary explanation that fustian is a matter of pompous banality. The disagreement comes in with respect to what is regarded as a case of fustian. What one takes to be a grand passage, another may dismiss as pompously banal. This disagreement is not about the meaning of a word, that is, it is not a “semantic” one, but is, instead, about how one understands and appreciates the character of the passage. What may not have been noticed is that very similar considerations can apply to our judgments of people. Where one may see, say, exemplary generosity, another may find only wasteful extravagance.”¹⁶¹

If disagreement on aesthetic attributions was a matter of ascribing different meanings to *predicates* for A, there would not be a genuine disagreement about possession of A by O. And neither would there be genuine disagreement if aesthetic attributions were merely about *private* objects of experience, in the sense that my judgment that O is elegant would be a judgment about an object that is inaccessible to other perceiving subjects that deny my ascription of elegance to O. The relativist cannot have his cake and eat it too, here; if we genuinely disagree about possession of A by O, there must be a shared understanding of what it is for O to have A; if there is no shared understanding, disagreement cannot be invoked in favor of aesthetic relativism. And recognition of this fact does not force one to incline towards any particular theory about the ultimate nature of aesthetic qualities.

¹⁶¹ Tilghman (2006: pp. 161-162).

The relevance of these observations for an aesthetic theory of music is that, on the one hand, to say that to experience sounds as musical is a special case of experiencing sounds aesthetically does not compel the endorsement of this or that particular theory of aesthetic properties; the theory requires only that there is such a thing as experiencing sounds aesthetically. On the other hand, it indicates that disagreement about the relative value of a musical piece is compatible with agreement on substantive aesthetic descriptions of that piece. For instance, one can note the *nobility* of a French horn melody in the last movement of Brahms's first symphony Op. 68, or the already mentioned *cragginess* of Beethoven's *Hammerklavier* sonata, while not sharing with others, who agree on experiencing the same features of the music, an assessment of the work's relative merit or while differing in personal preferences (e.g. one can perfectly perceive "romantic" music of the nineteenth century as having the relevant substantive aesthetic properties, while being more drawn to the music of other periods). Of course, this doesn't mean that people always agree on the attribution of substantive aesthetic qualities, and neither does it mean that disagreement on verdictive attributions will be radical, given agreement on substantive attributions. Where people agree on substantive aesthetic attributions they will tend to converge on verdictive attributions; but ascriptions of relative value are one thing, verdictive attributions of beauty and ugliness are a distinct matter. Disagreement on relative evaluations is compatible with agreement on verdictive attributions (being able to experience the beauty of Brahms' first symphony is compatible with different orderings of one's musical preferences). In other words, aesthetic impartiality is compatible with different sets of preferences.

Now, why do I find Zemach's conception of aesthetic qualities as desire mediated phenomenal properties useful? First of all, because I believe it makes less mysterious our ability to experience things as having aesthetic qualities. A degree of care is required not to identify squarely this idea with the denial of the Kantian thesis that *pleasure* in the beautiful is independent of desire (even though Zemach seems to incur in this, in his attack on the idea of "disinterestedness"). This Kantian thesis merely distinguishes the experience of pleasure in an object's beauty from the experience of feeling pleasure in something because it satisfies an occurrent desire of ours (e.g. the pleasure in seeing one's soccer team score a goal – an example of "interested pleasure"). That pleasure taken in a thing's aesthetic qualities is independent of occurrent desires is compatible with the idea that aesthetic qualities are conceptually connected with our nature as "creatures of desire" (even if Kant himself would reject this view); as historically and culturally situated cognitive agents with limitations and biases, that is, the kind of beings that experience some combinations of elements as more or less fitting our cognitive and perceptual abilities, and thus as more or less coherent, unified, rich, noisy, amenable, distressing, vibrant or peaceful.

Take the example of two visual patterns of lines and colors of which one is said to be more “coherent” than the other. This attribution is far from being literal, since a visual array of lines and colors, in purely physical terms, is as “coherent” as any other (coherence is predicated literally of assertions – the property in virtue of which the truth of different assertions is mutually compatible). As applied to visual arrangements, the predicate “coherent” is relative to what fits more or less nicely with our perceptual and cognitive abilities – some visual patterns are experienced by us as “hanging together”, others as “confused” or “noisy”, and the attribution also refers to how the different aesthetic qualities of the pattern combine with each other (delicacy taken by itself is a good-making quality, but in certain contexts it may be out of place – different aesthetic qualities may “clash”, while others “hang together”). Zemach frames this idea in terms of “cognitive desires”: “We see things via desire. Specific aesthetic properties are modifications of phenomenal properties by any manner of desire; the general ones are perceived when things are *seen as* suited to cognitive desires only. The latter are ways of sorting things by fitness to our cognitive apparatus, that is, by their degree of perceived cognitive goodness.”¹⁶²

Some examples lend more clarity to this notion of experiences tinted by cognitive desires. Take the case of sublimity or awesomeness. Things that are appropriately described as “sublime” are those things that greatly surpass our powers. I can experience the depiction of a mountain range as a representation of a sublime thing only if I’m aware of the relative dimensions of mountains and human beings. My pet cat looks *cute* to me partly due to my awareness of the relative dimensions of cats and human beings; if I try to imagine the cat as a creature several times my size, it may no longer look cute, but sinister or menacing. “To appreciate a tree in fall you need to know about death; to see the beauty of a deer you need to know what the function of eyes and legs is; to see the Alps as sublime you need to know how high most hills are.”¹⁶³ This has important consequences for how one is to conceive *formalism* about aesthetic qualities. If we take formalism to consist in the narrowing of experience, to the effect that appreciating something aesthetically would involve severing the connections between our immediate experience of objects and the rest of our experience (and also severing each aesthetic judgment apart from our other aesthetic valuations), it’s hard to make the idea of aesthetic appreciation intelligible under that notion. On the contrary, to experience things aesthetically involves experiencing them in terms of *something else*, namely, as standing in certain relations to *us*, as affecting us in determinate ways, but also as standing in certain relations to other things in the world. This makes aesthetic qualities something much more

¹⁶² Zemach (1997: p. 107).

¹⁶³ *Ibid.*, p. 36.

intelligible than to conceive them as that which remains *after* we sever all connections between our immediate experience of an object and the rest of our experiences and evaluations.

I said that to experience things aesthetically involves experiencing them “in terms of something else”. This rather vague notion was already partially explained in our brief digression through the workings of aesthetic language, that is, the deployment of words for aesthetic use (either these are primarily aesthetic predicates – predicates whose primary application is in the issuing of aesthetic descriptions; or primarily non-aesthetic predicates being put to aesthetic use). We saw that one conspicuous feature of the aesthetic use of words is that they are employed non-literally. When we describe a melody as “delicate” we are not ascribing to it a dispositional property that it would share with fragile things. Yet, such use of the word “delicate” retains a connection with its literal use; that is, we are able to think of the melody as “delicate” because, to an extent, we experience it “in terms of something else” – our experience of the melody exploits a mapping between the domain of sounds and the domain of other objects to which the term “delicate” applies in its primary, non-aesthetic use. What this means is that possession of the non-aesthetic concept expressed by the predicate in its primary use is connected with our ability to apply the aesthetic concept expressed by the same predicate when it is used non-literally. The same connection occurs in the previous example where the predicate “coherent” was applied to a visual pattern (our ability to experience the visual pattern as “coherent” is connected with our mastery of the non-aesthetic concept, as it describes a logical property of assertions). There are issues about the precise nature of this connection, but that need not worry us for present purposes. This doesn’t mean, of course, that *all* figurative use of a predicate constitutes an aesthetic use of that predicate. So, what is the *criterion* to distinguish aesthetic uses from non-aesthetic uses? Most likely, there is no way of specifying *a priori* a criterion for when a non-literal description of something counts as an aesthetic description. It will depend on the particular use of the description and it may vary with the context. Perhaps one way to shed some light on the matter is to circumscribe the aesthetic use of predicates to the domain of visual and aural appearances. (I leave open the issue of the aesthetics of other sense modalities, though it seems clear that attributions of delicacy to an object’s visual or aural appearance also involve some relation to tactile sensations.) This problem is connected, I believe, to the problem of characterizing the distinction between a *formalist* stance in aesthetics and a correspondent *anti-formalist* stance. This is a difficult distinction, to the extent that it is not clear how *formal properties* are to be distinguished from *non-formal properties*. One unsatisfactory way of characterizing formal properties is to say that the (aesthetic) formal properties of an object are those that are accessible in a direct perceptual experience of it. This is unsatisfactory because perceptual experiences of an object are affected by what one *knows* about the object, about the object’s relation

with other things in the world, etc., even though that knowledge is not itself part of the perceptual *content* of the experience. For instance, as well as seeing that this set of marks on a surface is a representation of a human being I also see that it is a representation of a *particular* human being (be it a real or fictional person). The experience is wholly perceptual (as is the experience of hearing a certain sound as a word token) but the qualitative aspects of the experience are affected by the knowledge that allows me to interpret the image as a representation of *that* particular person, even though this knowledge is not itself part of the perceptual content of the experience. Likewise, our perceptual experiences are “guided”, in a variety of ways, by our ability to make more or less fine-grained linguistic descriptions of what we experience. To use a very simple example, my experience of this object *as a chair* is wholly perceptual, though the concept CHAIR is not part of the perceptual content of that experience. The ability to subsume what one sees under sortal concepts modifies the experience (a perceptual experience is not merely a sum of brute stimuli; we perceive things as ordered or structured in a variety of ways). Also, our perceptual experiences are affected by differences in “background”: my knowledge of the “categories” to which artworks belong (Walton 1970) – e.g. painting, byzantine mosaic, bass-relief, etc. – enables me to distinguish what features are standard, contra-standard and variable in objects of that kind, which in turn will affect the way I experience the object (the idea being that if two indistinguishable objects are experienced as belonging to different art-categories, they will look different.) So, if what is accessible in perceptual experience varies with knowledge and background, it will not be helpful to specify “formal” properties in terms of what is accessible in direct perceptual experience.

Another unsatisfactory way of specifying these properties is via the contrast between *form* and *content*. This is because there is no clear boundary between what counts as form and what counts as content. One kind of properties that are traditionally considered as “content” properties are representational properties (*form*, by contrast, would be something like a “mode of presentation” of what is represented, a “structure” or a particular way of organizing the “content”). But representational properties often are relevant to determine a work’s “structure” to the point that it is not clear where “content” ends and “form” begins. (To take a well-known example from the literature about the intersection of ethics and aesthetics: in ancient Greek tragedies, the moral properties of the characters are crucial for the work to have the appropriate “structure” – e.g. the protagonist can’t be contemptible nor a moral saint.) And in non-representational artworks, as is the case with absolute music or non-figurative painting, the distinction breaks down, since there is no way to pull “content” and “form” apart. The “content” of a piece of absolute music is the sum of musical elements (notes, chords, themes, timbral texture, etc.) in certain relations with each other. But if rhythmically inflected sequences of notes and chords and timbres are the “content” of a piece

of absolute music, what exactly is its “form”? If we pick the same notes, chords, rhythms and timbres and rearrange them differently, the result will be another, distinct piece of music. And this is compatible with both arrangements, say, exemplifying something like sonata-form. But to specify the form of each particular arrangement we cannot abstract from the particular notes, chords, etc. that make it up.

A more promising way of characterizing the distinction between formalism and anti-formalism in aesthetics is perhaps in terms of *bare appearances* (the appearance of something independently of how it really is), as opposed to the appearance of something in connection with how it really is (this has been recently defended, e.g., by Nick Zangwill regarding the aesthetics of inorganic nature).¹⁶⁴ So, for instance, a formalist about the aesthetics of a particular range of objects will say that what matters for the aesthetic appreciation of those objects are their bare appearances. What this means is that in order to appreciate and enjoy the beauty and other aesthetic features of the relevant objects, one is not required to have knowledge of that object’s true nature, of how it really is. Furthermore, one can be a formalist about a certain range of objects without endorsing formalism across the board, for any category of objects. Neither is one compelled to be either thoroughly formalist or thoroughly anti-formalist. There is the intermediate stance that says that an object’s aesthetic qualities depend *both* on bare appearances and on the object’s true nature. This seems to be the sensible view regarding artifacts and *functional* things in general: where things have functions this is relevant to their aesthetic appreciation, since functional things may differ aesthetically in the way they perform their functions. For functional objects, their appearances may be more or less aesthetically apt *expressions* of their functions. For instance, the fact that a *prima facie* good-making aesthetic feature may be “out of place” in a particular kind of artifact is connected with the norms set by the object’s function. One example of relevant functional property is that of representational functions: a visual pattern may be aesthetically good *as* a representation of object O, without it being the case that it would be aesthetically good if considered as an abstract, non-representational pattern. Functional beauty is beauty that depends on the concept of a thing’s function or purpose. This is what Kant referred to as “dependent beauty”, by contrast with “free beauty”. Dependent beauty is beauty that depends on non-aesthetic functions.

This characterization, unlike the preceding ones, allows for a plausible distinction between the formalist and the anti-formalist stance. In particular, it does not require that whatever aesthetically relevant features that are perceptually accessible in the experience of an object must be conceptually

¹⁶⁴ See Zangwill (2013).

disconnected from the rest of one's experience nor that we must take the object in abstraction from its relations to other things in order to experience its aesthetic qualities. It merely prescribes that where an object's aesthetic qualities are *formal* that means that its bare appearance, regardless of how it really is, is what matters. Zangwill gives the example of the beauty of clouds, which seems to be tied to a perceptual illusion, the illusion of solidity. Clouds "appear to us as fluffy, bouncy, solid objects. Perhaps they look as if they are made of cotton-wool."¹⁶⁵ This illusion of solidity is crucial to our appreciation of their beauty. (Consider the aesthetic description of clouds as "magnificent massive floating castles".¹⁶⁶) And it is a further example of how the appreciation of a thing's aesthetic qualities involves experiencing them "in terms of something else", in this case, in terms of our previous grasp of the concept of solidity (as well as other concepts), as it applies straightforwardly to things that do exemplify the property of solidity. Now, an anti-formalist will want to defend the view that to appreciate things aesthetically we must know about their true nature, so that cases such as that evoked by Zangwill, that involve illusions such as the illusion of solidity are *inappropriate* ways of appreciating the aesthetic qualities of things. So the crucial concepts to ground the distinction between formalism and anti-formalism are those of *appearance*, *function* and a connection between appearance and knowledge of how things really are. Formal aesthetic properties are those aesthetic properties that depend solely on formal non-aesthetic properties (bare appearances); non-formal aesthetic properties are those that depend, at least in part, on broad, non-formal non-aesthetic properties (such as historical, functional and contextual properties). A moderate formalist will hold that for some classes of objects, at least, both bare appearances and knowledge of function are aesthetically relevant. This is *a fortiori* the case of artworks, which, as artifacts, are a subclass of functional entities. Concerning artworks, the moderate formalist accepts that both formal and non-formal properties can be aesthetically relevant. (We have now a more robust distinction between formal and non-formal properties, in terms of bare appearances and a more broad class of properties – including contextual properties, such as historical properties about a work's origin and the artistic category to which it belongs – on which the aesthetic character of an object may depend.) An anti-formalist about art, by contrast, will deny that *only* formal properties can ever be a proper basis for an aesthetic appreciation of artworks.

This raises some issues concerning absolute music: purely instrumental music without title, words or programme (or as Peter Kivy refers to it, "music alone"). Even the moderate formalist will be tempted to regard the relevant properties of absolute music as purely formal, that is, even the

¹⁶⁵ Zangwill (2013: p.585).

¹⁶⁶ *Ibid.*, p. 586.

moderate formalist about art will be tempted to endorse extreme formalism about absolute music. If one endorses extreme formalism about absolute music (or at least most absolute music), one commits to the view that the aesthetic appreciation of this kind of music does not differ from the aesthetic appreciation of inorganic nature. The aesthetics of pure music would be like the aesthetics of clouds. This point was forcefully put by Nick Zangwill in a recent article on the social dimension of music listening:

“Listening to music is an isolated and lonely encounter with another world, a disembodied world of beautiful sound, far from the world of human life. . . . To humanize music is to desecrate it. Music is inhuman, and awesome because of it, like stars in the night sky.”¹⁶⁷

As we’ve seen, however, such a position does not conflict with the view that pure or absolute music is often experienced “in terms of something else”, one of the most conspicuous examples being that of experiencing music in terms of emotions; not in terms of representation, arousal or expression of emotions, but in terms of the music itself exhibiting emotional qualities (properties which music can no more exemplify literally than clouds can literally exemplify solidity). Experience of (absolute) music in terms of emotion, for the formalist, is as much a *formal* property of music as the appearance of solidity in clouds. The anti-formalist about absolute music will hold that music’s emotional properties depend on some *literal* connection between the emotional qualities we detect in music and the actual emotional states of some agent – the composer, the performer or the listener. (One such position is that we *imagine* the music’s unfolding as the expression of a real or fictional *persona*.)

One advantage of the formalist view is theoretical simplicity: we don’t need a separate account of the emotional qualities of music, distinct from the account of music’s other expressive features. Some versions of anti-formalism about absolute music will explain the emotional expressive qualities of music by appealing to a literal connection between our experience of such features in the music and actual emotional states of cognitive agents, while explaining non-emotional expressiveness (qualities that lend expressiveness to the music but are not appropriately described with emotion predicates) in terms of what is perceived in the music itself. Examples of the properties I’m referring to are: powerfulness, vividness, sprightliness, sereneness, solemnity, somberness, eeriness, etc. To the formalist, both these and the emotional qualities of music are of one kind – they are aesthetic qualities of the music; qualities we refer to with predicates that for the

¹⁶⁷ Zangwill: (2012: p. 389).

most part do not apply literally to the music, since sounds do not literally exemplify the properties described by those predicates, analogously to the way clouds do not literally exemplify solidity. I will explore in more detail in the next section the phenomenon of “aesthetic metaphor” as applied to the experience of music.

To conclude the present section, I want to leave a remark about the version of aesthetic formalism about absolute music just sketched. As it happens, I am skeptical of Zangwill’s thoroughgoing formalism about absolute music. The reason I am skeptical is that if we are to endorse aesthetic *functionalism* about music, that is, if we are to explain music in terms of the intentional realization of aesthetic properties that depend on sounds, I don’t think we can abstract from the *origin* of music in the insights and intentions of human agents when it comes to musical experience and the *appreciation* of music’s qualities. Music is not like inorganic nature. We do not listen to a piece of absolute music the same way we attend to the sounds of the Aeolian harp or to any sequence of purely natural sounds, even if these sounds are, to some extent, *music-like*.

4 Hearing sounds as tones: toward a natural-kind theory of music

What in every other art is still description is in music already metaphor.

Eduard Hanslick¹⁶⁸

Sounds are uncrystallized tones, tones that have not yet been realized.

Viktor Zuckerkandl¹⁶⁹

Roger Scruton has developed a thorough and compelling account of what musical experience is, though he never ventured a definition of music. In fact, he has manifested doubts about the value or usefulness of such endeavor.¹⁷⁰ Yet his view is one onto which a plausible definition of music might be piggybacked. In fact, our hanslickian definition of music as “the *art of tonal sound*” seems to effortlessly emerge from what Scruton says about the way we hear a sequence of sounds *as music*. This situation perfectly illustrates the distinction advised by Stephen Davies, between *definition of art* and *theory of art* (a distinction which is made obvious by the fact that at least some anti-definitional theories of art are theories about the nature of art). It is not enough to produce a definition with no obviously damaging counterexamples. For one, the same definition might be compatible with more than one background theory (there is at least one definition I can think of which is compatible with *all* conceivable theories, though it is also most uninteresting: “art is what is not not-art”; or consider Stephen Davies’ observation that “music is not [merely] patterned sound, it is the bodying forth of sound through human action”¹⁷¹, which is not offered as a definition); furthermore, it is not realistic to expect a definition of a social/artifactual kind to be so neat that it meets no purported counterexamples or hard cases. But in order to see whether the definition can

¹⁶⁸ (1986: p. 30).

¹⁶⁹ (1956: p. 2).

¹⁷⁰ For instance, Scruton (2007b, p. 238): “In particular the constant wrangling over the definition of art, over whether found objects are works of art, whether John Cage’s 4’33” is a work of music — even over whether works of art are types or tokens, whether notation determines identity, whether arrangements are identical with or versions of the originals, whether copies are versions or new works of art, and so on and so on — all such questions seem to me to get nowhere. For they can all be answered as you will, without casting any light whatsoever on why works of art are significant to us, and what kind of significance they have.” But I have doubts about these reasons. Even if the definition of art did not cast light on why works of art are significant to us and what kind of significance they have, that would not show that pursuing a definition is futile or pointless. On the other hand, I’m not so sure that the issue of definition does not cast light on the issue of value or significance.

¹⁷¹ S. Davies (2012: p. 548).

handle these hard cases, the definition itself is not enough. What one has to say in order to adjust or fine-tune the theory to deal with hard cases is part of the background theory that supports the definition.

As it happens, what I intend is not so much to piggyback a definition onto Scruton's view, but to take from it what I think is indispensable to an account of music's nature, without committing myself to every aspect of the theory, namely, his anti-realism and the idea of *metaphorical perception* (a perceptual state where something is represented as falling under a concept that doesn't literally apply to it), an idea that is, to an extent, also taken up by Christopher Peacocke in a number of essays on musical expressiveness. The views of Scruton and Peacocke about music are keenly related, though distinct, in that Peacocke's account is concerned only with music's expressive qualities, while Scruton digs deeper, in an attempt to describe those acts of perception and understanding in which hearing a sequence of sounds as sounds becomes distinct from hearing a sequence of sounds as a *musical* sequence.

So what is it to hear sounds as music? What are the fundamental components of musical experience? Scruton believes, as I do, that to learn what music is we must look for that knowledge in the attentive description of musical experience itself. We will not gain knowledge of what music is simply by observing what regions of the brain "light up" when one listens to music, important though those observations may be. For instance, whatever we may infer from the fact that music and language share brain resources, it is certainly *not* that music is a language "of sorts", that it is relevantly like language in some crucial aspect (resemblance is cheap, so music will be language-like in many *non-relevant* ways), nor that a sound sequence becomes a musical sequence in virtue of the same kind of properties that make a sound sequence become a meaningful linguistic sequence, though it may strengthen the view that music is a prerogative of language speaking animals like us. (I venture that it will be difficult to find any two reasonably complex human activities that *don't* share brain systems.) Nor will the observation that some brain area associated with the recognition of emotions is active while listening to music tell us whether this is because music arouses emotions without in any way representing or expressing them; whether the emotions supposedly experienced are those of the music's author, of the listener, of some fictional persona one imagines in the music, and so on. It will not tell us whether the relation between music and the supposed emotions is a literal relation of representing or expressing or whether emotion terms are only metaphorically applied to music, in such a way that emotion predicates either do not ascribe any properties whatsoever to the music or the properties they ascribe are not literally emotional properties. It does not tell us whether some emotional responses to music are more or less appropriate to the music and why. Music may "tap into" the brain resources we use to detect

emotions in people's behavior; but this tells us only that our capacities and dispositions for emotion are somehow engaged. That we recognize emotions in music (and, accordingly, make use of our capacities to recognize emotions in other things) is not an answer to the puzzle but a restatement of what puzzled us in the first place: how can we hear emotions in music? How are emotional descriptions of music appropriate to it? It tells us nothing about the nature of music or how *it* relates to emotion, if at all.

But I don't wish to discuss the relative merits of empirical research on how music works at the level of the brain, not least because I do not have the competence for it. The point I'm making is a rather modest one: that some questions are such that they are not susceptible of being answered only through empirical means. We must think and evaluate hypotheses that cannot be proven nor disproven only empirically. This is not to say, of course, that empirical investigation has nothing valuable to tell us about music. That would be obviously false and plainly foolish. For instance, only empirical research can tell us whether emotional responses to music exhibit some sort of cross-cultural uniformity. But we can't discover empirically whether our emotional responses to music are responses to music's representational properties or to some other kind of property. Likewise, empirical investigation cannot tell us what music is. If we were to define music as "what activates and excites music-specific neural functions or circuits"¹⁷² (assuming there are such music-specific structures in the brain) this would be only as illuminating as our previous ability to distinguish music from what is not music and no more illuminating than a definition according to which music is that which is made using musical instruments. Another difficulty is posed by the fact that music relies on many abilities that are not music-specific (e.g. "brain systems" that were not evolutionarily selected for music), so that to define music as that which engages those abilities would not take us very far.

¹⁷² In this paper, S. Davies (2012: p. 541) discusses several approaches to music's definition, one of which he terms "operational", and characterizes as an extensional definition, not as an attempt to uncover "the essence of the concept", as Davies calls it.

4.1 The components of musical experience: acousmatic listening

The first key component in hearing sounds as music, according to Scruton (1997; 2009), is what he terms the *acousmatic experience* of sounds. This is a concept he adapts from Pierre Schaeffer, the pioneer of *musique concrète*, who introduced the concept in the 1950's to describe the kind of experience invited by his own work with the medium of recorded sound. The concept of acousmatic experience refers to the perception of sounds in abstraction of their physical causes (the sound sources). Recording technology allowed this kind of detached experience of raw sounds, natural or artificial, and their further manipulation until the connection between them and the sound sources is completely opaque to the listener. The term *musique concrète* alludes to the direct manipulation of sound by the composer, which is then played back, not performed, as opposed to the “abstract” sound structures of traditional music.

The term “acousmatic” is derived from the Greek word *akousmatikoi*, meaning “one willing to hear”, which was used to refer those acolytes of the Pythagorean School that attended lectures where the lecturer spoke from behind a curtain so that the attendants would focus exclusively on his words. Since Schaeffer had already integrated the term in musical usage, it is a convenient one, though as good as any other. Whatever term one uses, what is relevant is the distinction it marks: sounds heard acousmatically are heard as events separate from the events that cause them and thus differ from how sounds are experienced otherwise, i.e., as acoustic cues allowing us to navigate the environment; in this way we hear the sound sources (the sound producing events) through the sounds they produce, and that is the focus of our attention: the information we can gather, via the sounds, about the sound sources. We may call this the *acoustical* experience of sounds: the way they are regularly attended as features of an environment, such that our interest in the sounds is not detachable from our interest in the immediate environment.

The acousmatic experience of sounds, on the other hand, occurs when sounds are detached, in our experience of them, from their sources; when we are no longer focused on sounds as features of their sources or of the immediate environment but as separate events, such that the focus of our attention are their properties *as properties of sounds*, exhibiting an order apart from the everyday physical world. The most obvious manifestation of this ability is given in the way we spontaneously individuate sound *sequences*, distinguishing them from other co-occurring “sonic individuals”. If we hear, say, the steps of a C major scale interspersed with noises, we will perceive the scalar steps as parts of a single sonic individual, which is partly obscured by the intervening sounds. Yet, from the point of view of pure sound, these are all indistinct elements of a single acoustic event, a chaotic clashing of sound waves, mindless patterns of vibrating air molecules. It is *we* who seek and

perceive *order* in these events, by having experiences where certain elements are foregrounded while others are made secondary or even ignored. Nature itself makes neither distinction nor privilege between sounds with regular or irregular wave patterns; it has no special interest in the relative prominence of certain overtones, responsible for the distinctive timbral profile of individual musical instruments; it doesn't care whether a musical line goes "up" or "down", or how a sequence of chords ends or "resolves" or "suspends in mid-air". It is in this sense that music happens "not in the world but in the theatre of the human mind."¹⁷³ Of course, the "theatre of the human mind" is itself a part of the world, so that music happens in the world after all. This is not to endorse Collingwood's conception of artworks as "mental entities" but to stress the fact that human experience is one of the *relata* in the relations that constitute certain objects as artworks, and certain sound events as *musical* sound events. Likewise, we don't need to endorse the view that sentences are mental items, severed from the objective world, to assert that in a world without thinking subjects no sounds, however intricately constituted, count as linguistic utterances.

While the pioneers of *musique concrète* wanted to make us attend acousmatically to everyday noises, which takes considerable effort when the sound sources are easily identifiable, as in some of the earlier pieces by Schaeffer (e.g., a train arriving at the station), the most common and spontaneous form of acousmatic experience is found in musical listening of the traditional kind (music made of melodies, rhythms and harmonies, rather than the clicking sound of a hundred metronomes).

"Music makes use of a particular kind of sound: an acousmatic event, which is heard 'apart from' the everyday physical world, and recognized as the instance of a type. This isolation of the pure sound event leads to a peculiar experience, which I have called the experience of tone. No longer does it seem as though the middle C that sounds is caused by someone blowing on the clarinet. Instead we hear it as a response to the B that preceded it, and as though calling in turn for the E that follows. . . . When we hear music, we do not hear sound only; we hear something in the sound, something which moves with a force of its own. This intentional object of the musical perception is what I refer to by the word 'tone'."¹⁷⁴

To understand what is meant by "hearing sounds as tones" we must take notice of the peculiar technical use that is being imposed on the word "tone", which is of itself an ambiguous word in musical usage, its meaning ranging from "sound of definite pitch" to any of the old "ecclesiastical

¹⁷³ Dutton (2009: p. 103).

¹⁷⁴ Scruton (1997: p. 19).

modes” (a kind of musical scale) and is even used as a synonym for “timbre” (as in “tone color”). Once again, what is relevant is not the word used but the distinction intended by its use. The experience of sounds as “tones” is one kind of acousmatic experience of sound, which is defined as the experience of sounds in abstraction from their sources. The sounds of *musique concrète* are intended for another kind of acousmatic experience, one directed at the features of sounds manipulated through recording technology. Sounds heard as meaningful utterances from a speaker constitute yet another kind of acousmatic experience, one in which an organization imposed on the sounds is made more prominent than the organization imprinted on them by the literal causality of sound production. The sound of the speaking voice is no more “environmental” than the rhythms and harmonies of a musical line or the “tweaked” sounds of a non-tonal electronic composition. Significantly, both “tonal” and linguistic sounds can be the object of a further “abstract listening”, as these were also used as compositional elements in works of *musique concrète* and later “avant-garde” experiments influenced by it.

The experience of sounds as “tones” doesn’t cover only those sounds whose organization includes the parameter of *pitch* (though that is roughly the standard use of the term). It includes, first and foremost, the experience of *rhythm*, without which tonal organization could not be perceived as such. It is through rhythm that a sounding C note can be experienced as “answering” the B that preceded it and “calling for” the E that follows. Non-pitched rhythmic sounds are heard as endowed with a life and movement of their own; in fact, no sequence of sounds can have melodic identity if it lacks rhythmic identity. Rhythm is what “energizes” a musical line; it is the feature in virtue of which music “moves” at all. Despite the fact that it is often so defined, it is not to be identified with mere regular, patterned or metric sound, neither is it reducible to a pulse, since all of these can exist in the absence of rhythm. Rather, *beat*, *pulse* and *meter* are *constituents* of rhythm, and some music is not even built on the divisive meter that is typical of western music. Often the beat will be merely implied by the melodic or harmonic texture, and some music has free rhythm (like medieval plainchant). Rhythm is not an intrinsic property of sounds, like frequency; rhythm is something that is heard *in* sound sequences by creatures endowed with the right cognitive and perceptual apparatus. It is a known fact that in listening to the clicking of a metronome we spontaneously group the sounds in pairs, hearing one of them as stressed – a feature that is probably derived from our constitution as bipeds. Hearing rhythm is not simply a case of hearing *organized sound*. Sounds can be heard as organized without being heard as *rhythmically organized*. As

Scruton puts it, “in hearing rhythms, we are hearing a kind of animation.”¹⁷⁵ Sounds perceived rhythmically are, in a way, perceived anthropomorphically, “as if they were living beings whose physiognomy is akin to ours.”¹⁷⁶

Rhythm is one of the parameters of this kind of experience of sounds. It has a sort of priority over other parameters in the sense that a purely rhythmical line, constituted by sounds of indefinite pitch, already has *musical* identity; a rhythmic sequence can “stand on its own”, so to speak. The other parameters are pitch, timbre (which Scruton excludes from the acousmatic experience while others, like Julian Dodd, include), melody, and harmony. The latter two are also dependent on rhythm, in that experience of melodic movement is not detachable from the rhythmic movement that underlies it, nor is the experience of harmonic progression conceivable in a sequence which has no rhythmic character. In such a sequence, one could still hear sounds as chords, as in Levinson’s imagined series of “colorful instantaneous combinations of chords of vanishingly brief duration”, or in Yves Klein’s *Monotone Silence Symphony*, which consists of a D major chord held for about seven minutes, followed by forty-four minutes in which the players are silent, but to hear harmonic *development* rhythm is a required element. I don’t wish to quibble with the possibility that pieces made from pure timbral variations of isolated pitches scattered through long periods of silence can’t be properly called “musical”.

There is a way of illuminating the acousmatic experience of sounds as tones by analogy with the experience of pictorial depiction, i.e., with the experience of seeing the depicted object *in* a set of marks on a surface. Experiences of pictures that depict their subjects are said to be twofold: they involve awareness of the picture as a set of marks on a surface *and* awareness of the picture *as of* a depicted scene or object. Now, imagine someone who is able to see the marks on the surface but not the depicted scene or object. One way to achieve this effect ourselves would perhaps be by squinting at the picture, or by isolating and magnifying a detail of it, to the point that it becomes an amorphous blob. Suppose that someone visually experiences the whole picture in the same way that we would experience the blob. That person would have a visual experience of the picture as a set of marks on a surface, but not as of a depicted scene or object. This experience exactly parallels the experience of hearing sounds without hearing an “order of tones” in them. It is what actually happens to people with the condition known as “amusia” or tone-deafness.¹⁷⁷ Arguably, sequences

¹⁷⁵ Scruton (1997: p. 30).

¹⁷⁶ Zemach (1997: p.105): In this passage, Zemach is not talking about sounds in particular, but of aesthetic qualities of things in general.

¹⁷⁷ See Oliver Sacks 2007, chapter 8. Steven Mithen, 2004, chapter 4.

of “toneless” sounds of differing pitches and durations are what our pets experience when we turn the stereo to our most beloved piece of music.

Up to now we have identified two major components of the experience of music. The first is the acousmatic experience of sound, which music shares with other sound-oriented activities. The second is the experience of hearing sounds as “tones”, which is one of the modes of acousmatic experience. This second component, in its turn, has its own set of elements: rhythm, pitch, timbre, melody and harmony. We have seen that rhythm has a sort of priority over the other elements, in that it can stand on its own (though it is perfectly conceivable that a piece properly called “music” would lack it). None of the terms with which I described the elements or “parameters” of musical experience refers to mind-independent or purely acoustical properties of sounds. Frequency is one such property: independently of how they are perceived, all sounds are patterns of vibration imprinted on the air (or another medium) by a vibrating sound source. But pitch is not the same as frequency, though it depends on it. Sounds perceived as of the same pitch can vary in frequency without that variation being registered at all by perceivers; some intermediate frequencies are not perceived as separate pitches at all (as having an “identity of their own”), but as “out of tune” versions of other pitches in a musical scale (this is why some experiments in microtonal music sound simply like music played out of tune). The human ear has its own biases, which are biologically predetermined – like the phenomenon of octave equivalence – which partly explains why, of all the immense varieties of ways that the pitch continuum can be divided¹⁷⁸, there is a cross-cultural predominance of pentatonic and heptatonic scales¹⁷⁹ (though differing widely among themselves), of which the former are practically a musical universal (though, once again, there are many different kinds of pentatonic scales possible). Dividing the pitch continuum allows us to have musical scales, which are particular ways of structuring the sounds contained within the interval of the octave. Scales are the raw material from which tunes and melodies are made, and different scales generate a different kind of “movement” that the musical line is capable of, with its own distinct “character” or “mood”. For instance, some music that sounds peculiarly “Spanish” is made using scales built on the Phrygian mode (one of the old ecclesiastical modes prevalent in western music before the advent of major/minor diatonic tonality). Analogously, tunes made from a scale containing the interval of “augmented second” (such as the harmonic minor scale or the double

¹⁷⁸ If the layout of our keyboard instruments represented the natural succession of the overtones series, without imposing a “human organization” onto that continuum, there would have to be 77 keys between any two notes an octave apart. As it happens, we have settled for 12 notes, allowing us 17 “keys” or “tonalities”, of which 5 are “enharmonic”, that is, sonically indistinguishable, at least in the system of equal temperament.

¹⁷⁹ That is, scales made of five and seven tones, respectively.

harmonic scale) will have a distinct “eastern” flavor. Of course the music itself has no “ethnic essence”, and our tendency to think otherwise will be partly due to the contingent fact that some scales or modes kept being used in some places while being abandoned in others. Nor is it the case that all music written in the Phrygian mode will have a “Spanish flavor” – take for instance Vaughan Williams’ well-known *Fantasia on a Theme of Thomas Tallis*, which is supposedly “quintessentially English”. Music made everywhere in the world has the same dependence base: the nature-given pitch continuum, the overtone series, and the universal biases of the human ear, favoring musical movement generated through scales of five and seven tones (but occasionally also favoring inventive departure from those templates, not to mention the wide variety allowed within each “template”).

I said before that Scruton excludes timbre from what we are calling “the experience of sounds as tones”, though others (e.g. Julian Dodd, Andy Hamilton)¹⁸⁰ include it (and my allegiance falls with the latter on this topic). Julian Dodd (2010: p. 36) has named the view that includes timbre among the elements of musical experience “modified acousmaticism”, and I follow that terminology here. But why does Scruton exclude timbre from the constitutive elements of “experiencing sounds as tones”? To understand this, one must bear in mind that Scruton’s theory is of an anti-realist bent. In his view, “tone” is the “intentional object” of musical experience, “one that we hear in sounds, but which has properties that no sound can have: for example direction, energy, a kind of internal ‘wanting’, together with relations of attraction and repulsion towards other tones.”¹⁸¹

“Tonal organization” is not “out there” in the sounds.” Tones” are rather features that we hear *in* the sounds, analogously to the visual experience of seeing a dog *in* a picture, which is literally a set of marks on a surface. This account of tone parallels Scruton’s view on the experience of aesthetic properties as a kind of “aspect perception”, the same kind of experience we have when seeing faces in clouds or a duck in the ambiguous duck-rabbit figure. He follows Kant in that for him “beautiful” and all the other aesthetic predicates do not refer actual properties of the objects but only ways in which we experience objects. For him, aesthetic judgments do not ascribe properties to objects, they do not express *beliefs* and the aesthetic concepts deployed in such judgments are not *asserted* of those objects, in a similar way that “this is a dog” is not asserted of the marked surface in which we see the dog picture. For Scruton, the experience of sounds as tones, as the aesthetic experience of

¹⁸⁰ Composer Aaron Copland, in his book *What to Listen for in Music* (2002), also includes timbre or “tone color” among the “elements of music” or, more precisely, of the musical experience of sound.

¹⁸¹ Scruton (2004: p. 186).

objects in general for Kant, is a kind of *imaginative experience*.¹⁸² We can now begin to understand why timbre would be excluded, in such a view, from the constitutive elements of musical experience. The idea is not, obviously, that in experiencing sounds as music we do not experience timbre. Rather, the point is that timbre is not among those qualities that are responsible for the musical nature of the experience. For instance, we experience sung words as music; yet, the words and their meanings are not part of what makes the experienced sound event a *musical* event. For Scruton, it is the organization of sounds in terms of pitch, rhythm, melody and harmony (one or more of which may be absent) that defines the experience of the sounds as *musical*. An implicit assumption so far is that musical experience is primarily experience of “tonal” organization in sounds (in our technical sense of “tonal”) and that other kinds of sound organization called “musical” must be so named in a metaphorical or at least in a secondary, derivative sense – an extension of the term in its primary use, unintelligible if not for that primary use.

Going back to timbre, Scruton rejects it as a defining element of musical experience because he understands such experience to be one of the modes in which sound is experienced acousmatically (where “acousmatic” contrasts with the acoustic experience of sounds – sounds experienced as revelatory of their causes; as acoustic cues to navigate the environment, not as separate events with “a life of their own”). So, if music crucially depends on acousmatic experience, then nothing that is not an element of the acousmatic experience of sounds can be an essential constituent of musical experience. Since he places timbre together with other intrinsic acoustic features of sounds, like frequency, he makes no room for timbre among the features in virtue of which an experience of sounds is a musical experience. Timbre would be one of those features of sounds that are experienced whether or not the experience is a musical one.

This view can be challenged though, and I’ve already hinted at that challenge a while back, when I said that mindless nature makes no discrimination between patterns of relative prominence in overtones (the feature responsible for our hearing a note in the clarinet as different from a note of the same pitch, intensity and duration in the oboe). We saw that pitch is a relational feature of sound, which Scruton does not call into question. I argue, likewise, that timbre has a relevantly similar relational component. We do not experience the peculiar “tone color” of the oboe or the clarinet as a matter of the relative prominence of overtones, just as we don’t perceive pitch as a matter of frequency (we don’t perceive frequencies, though we measure them) nor color as a matter

¹⁸² “I refuse to define beauty as a property of something; rather, I define the interest in beauty as a state of mind of the person observing something.”(From an interview to the *Postgraduate Journal of Aesthetics*, conducted by Sarah Hegenbarth and Vid Simonti, 2011.)

of wavelength (that is not the content of our color experiences). If we had to justify our practice of building oboes and clarinets, we would not invoke differences in the relative weight of overtones as a reason, because we have no reason to strive for the making of an instrument that issues sounds with those features except for the fact that those features sustain *other* features that we are interested in. Namely, the distinct tone color of the oboe has a unique *aesthetic profile* (it is sometimes said, e.g., that the oboe has a more “nostalgic”, “poignant” sound, in contrast to the more “joyful” or “playful” sound of the clarinet, though such descriptions will not be true in every circumstances), in virtue of which certain melodies in a musical piece are more suitably played with an oboe than with any other instrument. So, to exclude timbre from the inherently musical features of the musical experience is akin to exclude color from the organizational aspects of a painting.¹⁸³ When we speak of tone color or timbre, our interests are not focused on purely acoustic properties of sounds, but in one specific aspect of their aesthetic profile. This is evident in the way we tend to describe timbral qualities in metaphorical terms that function as aesthetic predicates: a timbre is said to be “metallic”, “mellow”, “sweet”, “warm”, “mournful”, “grating”, “shrill”, etc., yet these also refer to “properties that no sound can [literally] have”. Of course, non-musical sounds can have aesthetic properties without thereby becoming musical, so that even if some timbral properties are also aesthetic properties of the sounds, that doesn’t show timbre to be one of the constitutive elements of music. However, my reply to this is the following: in realizing particular rhythms, melodic lines or harmonies, the composer is trying to bring into fruition a particular, even if not exhaustively describable, set of aesthetic properties. But that particular aesthetic profile he has in mind doesn’t depend only on “pitch structure” and rhythmic sequence; it also depends on timbral aspects, such that the same melody will instantiate a very different aesthetic profile when realized in different instruments of widely varying tone colors. The reason we are interested in the realization of pitch structures with a particular rhythmic character is that we are interested in the aesthetic result (once again, purely acoustic properties of sounds will not figure in a justification of our interest in music). Our interest in music is not a Pythagorean mystical fixation in the acoustic instantiation of certain

¹⁸³ For a statement of this objection against (non-timbral) acousmaticism, see S. Davies (2008; 2011, Ch. 11). In this paper, Davies argues for “instrumentalism”, a view in the ontology of music according to which for a sound event S to be a performance of work W, which was, say, written for the clarinet, it is not sufficient that the sound event be produced in accordance with the composer’s specifications regarding pitch, rhythm and timbre, but it requires also that the performance is carried out with the same (kind of) musical instrument specified. In other words, a Perfect Timbral Synthesizer won’t be enough to produce a genuine performance of W. “Timbral Sonicism” (the view he argues against) is the view according to which while timbral properties are work-determinative, it is not necessary that they be produced using instruments of the kind specified by the composer. “Pure Sonicism” is the view that timbral properties are not work-determinative. This view was once defended by Peter Kivy (1993: pp. 75-94). Davies’ objection can thus be used by the defender of modified acousmaticism against the exclusion of timbre from the set of music-relevant properties, even if we stop short of accepting instrumental specifications as work-determinative.

“good” ratios. Music is about our aesthetic interest in sounds, not about venting in sounds our aesthetic preference for certain numeric relations. But if “sound structures” are to have any determinate aesthetic profile, choices regarding timbre have to be made, with differing aesthetic results. This is true even in a cultural context where the conventions regarding the specification of instruments are loose, such as it was with “pre-modern” instrumental music. The fact that choices of instrumentation were left to the performers makes timbre a no less important factor in how the “sound structure” is to be realized.

Though we usually identify timbral properties of sound by reference to sound sources (musical instruments), this is not a necessary constraint on perceiving the timbral qualities themselves. It is often the case (even more so with the availability of recorded performances) that we become acquainted with a particular timbre before we come to know the instrument that produces it. I remember hearing the *duduk*, an Armenian double reed wind instrument, for the first time in Peter Gabriel’s soundtrack for Martin Scorsese’s film adaptation of Nikos Kazantzakis’ novel *The Last Temptation of Christ*.¹⁸⁴ I only learned about the instrument itself years later, though I never failed to recognize its distinctive “mournful” timbre in other occasions. Should that sound have been generated for the first time in a Perfect Timbral Synthesizer (henceforth, PTS), an imaginary electronic device capable of generating “tone colors” in such a way that we couldn’t tell, by ear alone, whether they were electronically or mechanically generated, it would be no less distinctive; though there’s room to argue that its aesthetic properties are to an extent affected by our knowledge that it issues from a particular instrument, rooted in a particular musical tradition; in this case, that of Armenian folk music. But this is not essentially different from a case where the aesthetic properties of a melody are affected by our identifying their contour as typical of a particular musical tradition.

Furthermore, all experience of “tones” is timbrally inflected. Even an electronically generated “pure sine wave”, with no pattern of clashing overtones, has a distinctive profile derived precisely from that difference from all other tonal sounds with which we are acquainted. And contrasts of tone color are crucial for the organization of “musical space”, the same way that colors in a picture have important structural and symbolic functions. (Our experience of Titian’s painting *Bacchus and Ariadne* changes subtly once we learn about the restrictions of the time imposed on the use of ultramarine blue.) Suppose that we produce a replica of a known painting where each color is

¹⁸⁴ For a good display of the qualities of this instrument (as well as others), see the recording by Jordi Savall & Hesperion XXI, *L’Esprit d’Armenie*, AliaVox, AVSA 9892.

replaced by its complementary color (an “inverted spectrum” version of the painting, so to speak). Such change would result in drastic aesthetic differences from the original picture.

As an example of timbre contributing to the organization of musical space, one could suggest Webern’s 1935 orchestration of Bach’s six-part *Ricercar*, from *The Musical Offering* (BWV 1079), an example Scruton himself, oddly enough, uses to illustrate a point about melody:

“Even the most powerful musical Gestalt can be sent into the background and replaced by another. Consider Webern’s orchestration of Bach’s six-part ‘Ricercar’ from the Musical Offering: here the orchestration compels you to hear Bach’s melodic line as background, the foreground being occupied by the short motifs which, for Webern, form the true substance of this extraordinary work”¹⁸⁵

The effect described in this passage is achieved, precisely, through timbre, using the technique described by Schoenberg as *Klangfarbenmelodie* (“tone color melody”), which consists in breaking the melodic line, passing it through several instruments of distinct tone color. Later in the text and commenting on the same musical example, Scruton even comes close to recognize the role of timbre in structuring musical space:

“in [Webern’s orchestration of Bach] the melodic line is broken into motifs, and stuttered out in timbres so opposed that the piece seems as though pulverized and reconstituted out of tones that Bach would never have imagined.”¹⁸⁶

The acousmatic experience of sound consists simply in experiencing the sounds in abstraction from their physical causes. By itself, this feature does not yield us musical experience, musical experience being rather one variety (among others) of acousmatic experience. Acousmaticism is a thesis about musical experience, and about that kind of experience acousmaticism says that

“In order to listen with understanding to a sequence of sounds *as music*, we need only attend to the qualitative nature of the sounds themselves; we do not need to identify their causes.”¹⁸⁷

¹⁸⁵ Scruton (1997: p. 46).

¹⁸⁶ *Ibid.*, p. 99.

¹⁸⁷ Dodd (2010: pp. 36-37).

In other words, for someone who holds the acousmatic thesis, the music-relevant features of a stretch of sound are certain *phenomenal* features of it whose understanding as such doesn't require us to go beyond the sounds, delving into their physical causes, the sound generators and the particular actions that go into producing the sound (which doesn't mean we don't have to think about the sound stretch as a product of human agency, causally downstream from certain insights and intentions). We may be completely ignorant about these elements without thereby being prevented from listening *with understanding* to the stretch of sound as a stretch of music. Our understanding may be widened by awareness of how the sounds are produced, e.g., how a certain expressive feature is achieved in particular instruments, but that portion of our understanding thus widened is not, strictly speaking, *musical* understanding (though it is *about* music). For acousmaticism, musical understanding is completely inherent to the phenomenology of sound; it is an understanding that shows itself in more or less appropriate descriptions of the music in terms of "emotion, motion and height", that is, in (typically) metaphorical descriptions of music with more or less appropriate aesthetic predicates. Even a purely technical ("cogs and wheels") description of music using nothing but technical jargon (which is a difficult thing to achieve, given that the technical jargon itself is riddled with metaphors of "emotion, motion and height") will only have a point given the availability of an aesthetic description of the music (in terms of "emotion, motion and height"). In other words, the technical descriptions make sense because what they describe sustains the phenomenal properties of music which are the focus of our interest as listeners. In fact, "technical" descriptions of music are only *musical* to the extent that there is a specific kind of aesthetic phenomenology of sound events, the one we are describing under the head of "hearing sounds as tones". Eliminate that phenomenal character and there is nothing "musical" left about terms such as "dominant seventh chord" or "plagal cadence". Our interest in plagal cadences, for instance, is not explained in terms of "moving from the subdominant chord to the tonic at the close of a harmonic progression". Once again, no such description would figure in a justification of our interest in plagal cadences. Rather, our interest in plagal cadences, *that is*, our interest in moving from the subdominant chord to the tonic, is explained in terms of the specific *phenomenal* character of that cadence (which is the typical cadence heard in the "amen" ending of old church hymns; for which it is also known as the "amen cadence"). In other words, "technical" descriptions are just one way of describing the focus of our interest in music (they afford us a kind of blueprint to produce sound events with a relevantly similar aesthetic profile); it is not that in virtue of which the focus of interest *is* the focus of interest.

The opponent of the acousmatic thesis may argue that in many circumstances our understanding of a musical sound event will be impoverished if we ignore specific features of the *performance*.

(Andy Hamilton, for instance, argues that musical experience is *twofold*, in the sense that it includes both acousmatic and non-acousmatic features of sounds.¹⁸⁸) This will be particularly acute in musical traditions for which performance is the focus of attention. In fact, western “classical” music has only become more or less paradigmatic of acousmatic listening in virtue of a) the prominence within it of purely instrumental (or “absolute”) music since the eighteenth century, and b) the recording technology that allows us to have musical experiences in the absence of an actual experience of the performing musicians. In his recent book, *Understanding Music*, Roger Scruton construed such objections by way of an analogy with painting. In this way, the opponent of the acousmatic thesis will argue that it is no more correct to say that musical understanding is completely inherent in the phenomenology of sound, without reference to how it is physically “realized” or “embodied” than it is to say that a figurative painting by, say, Van Gogh can be understood merely as a figurative image, without reference to features such as the brush strokes that imprint the “visible residue of human action” on the painting’s surface.

Scruton’s answer to this challenge consists in comparing the experiences of someone who only sees brush strokes but no image (no depictive content), and of someone who sees the image but no brush strokes. The first is not experiencing the painting with understanding (no matter how much “aesthetic attitude” he invests on his act of seeing), in fact, he is not experiencing *the painting* at all, whereas the second is having the experience of “seeing figuratively” and so has the kind of understanding appropriate to the experience of the painting, even if he also misses something. The same applies to the person who experiences a stretch of sound as music, while completely ignorant of how the sounds are produced by musicians in performance.

The answer I’d favor is somewhat different, though I basically agree with Scruton’s verdict here. It is tempting to establish a parallel between the brush strokes in the above example and timbral properties in our musical examples. So it seems that, given the aptness of the analogy, Scruton proves his point about timbre with this example of figurative seeing without the awareness of brush strokes. Yet the brush strokes have a crucial *structural* role in figurative seeing. Namely, there are at least two ways of seeing the figurative image which are equally compatible with the “physical realization” in the marked surface. Suppose the image depicts the face of a human being. Depending on how one understands the brush strokes, these could make the difference between seeing the image as a) a particularly expressive depiction of a normal looking human face or b) a

¹⁸⁸ See Hamilton (2007, chapter 4).

naturalistic representation of a hideously disfigured human face.¹⁸⁹ And there is no figurative seeing of the depicted human face that is *not* “inflected”, either in a naturalistic or a “non-naturalistically expressive” way. Furthermore, the expressive character imprinted on the image by the brush strokes is to some extent independent of our awareness of how that stylistic feature of the image was achieved, analogously to the way the “mournful” timbre of the *duduk* is detachable from an awareness of the mechanical means by which that aesthetic/timbral profile is achieved. There is also an analogy here to be made in the way that figurative seeing is “inflected” by the brush stroke component, and the way the same melody is inflected by being diversely harmonized, so that if harmonic inflection is a music-relevant feature, then timbral inflection is no less relevant, since it can change the character of the melody in no less drastic ways *and* given that timbral qualities are no less *phenomenal* than pitch and rhythm.¹⁹⁰ “Mournful” is a better description of the *duduk*’s timbral profile than a specification of the relative weight of overtones in a sound issued from that instrument, a feature on which the “mournful” quality of the sound depends. The same phenomenal profile would be achieved if there were no such musical instruments and the relevant timbral properties were produced by a PTS if one were available. So, there is a relevant difference between *musical* timbre and distinguishing sounds made by clashing metal or wood, in that the former consists of a distinct *aesthetic* quality of the instrument’s “voice”, one that we feel compelled to describe with aesthetic predicates such as “mournful”. To hear the *duduk*’s timbre is also to hear the mournfulness (or any other of its aesthetic features, described with a phenomenally appropriate predicate) *in* its sound.

So much for timbral properties. The other two (sufficient but not individually necessary) music-constitutive features of acousmatic experience are melody and harmony. Like pitch, rhythm and timbre, melody and harmony are features that we hear *in* the sounds. “Melody” should be taken here to include the distinct phenomenon of “melisma” – such as can be heard in the extended vocalizations on a single syllable of medieval plainchant. To the extent that hearing melody consists in hearing a kind of movement, melody is (as already noted) inseparable from rhythm. No sound

¹⁸⁹ I took this example from John Dilworth (2008: p. 351).

¹⁹⁰ A. Copland (2011 [elect. ed.]: ch. 7, para. 6): “At times, of course, a composer conceives a theme and its tonal investiture instantaneously. There are outstanding examples of that in music. One that is often quoted is the flute solo at the beginning of *L’après midi d’un faune* (*Afternoon of a Faun*). That same theme, played by any other instrument than the flute, would induce a very different emotional feeling. It is impossible to imagine Debussy’s conceiving the theme first and then later deciding to orchestrate it for a flute. The two must have been conceived simultaneously.” What Copland here calls “emotional feeling” is what I call the “aesthetic profile” of the melody above, which includes more than emotion-expressive qualities. About this example, the defender of timbral sonicism will say that the flute theme had to be conceived for the timbral profile of the flute, with no essential reference to the instrument itself, with which that timbral profile is obtained. The “instrumentalist” (which is also not a friend of acousmaticism) will beg to differ.

sequence can be a melody if it has no rhythmic identity. Melodic organization forms one added layer, over a rhythmic base. Harmony will provide yet another layer, and both will inflect the underlying movement (I'm thinking of the phenomenon called "harmonic rhythm"). Melody occurs when the underlying rhythmic movement sustains another kind of movement, that of *tonal* movement. While rhythm gives us an experience of "moving forward" in sound, melody adds another dimension, that of moving "up" and "down" in a virtual space (some writers on the subject, e.g. Malcolm Budd, will deny that this spatial metaphor is essential to the description of the phenomenon, but using it will make us beg no questions, since no one denies that the metaphor *can* be used to describe it; the dispute is rather about the spatial metaphor being eliminable or not.) In other words, in melody we experience sounds *as if* something was moving *through* them, across regions of the pitch continuum, with different organizations of that continuum (musical scales) generating different kinds of movement. For instance, the so-called "well-temperament" and our modern equal temperament allow the music to move freely across the "keys" of major/minor tonality (a phenomenon called "modulation"), without making the music sound "out of tune" (a melody or an entire piece of music don't have to stay on the same key from beginning to end), which is a wholly different kind of movement from the one we experience with the system of just intonation. The emotional "Sturm und Drang" of nineteenth century romantic music, not to mention sonata-form, and with it symphonies, string quartets and concerti as we know them, are unthinkable in the absence of such devices. This is because the ability to modulate between keys is (partly) what accounts for the "emotional instability" that we experience in such music: it allows music to travel away from the "home" key to ever more "distant" keys, from which it can either "return home", producing a sense of "resolution" and "appeasement" or "conciliation" (the recourse to metaphor is irresistible here), or it can remain in a state of "irresolution", producing an experience of "uneasiness" or "unrest". (This is of course, only a dim description of a small portion of the devices involved in this kind of music.)

Recall a previous quote from Scruton, describing the acousmatic experience of a musical line as that of a C note arising as an "answer" to the B that preceded it, and an anticipation of the E that follows it; as opposed to the literal order of causation, which flows from sound source to sound event, in the absence of "real" movement (the sounds don't literally "go" anywhere). This "virtual movement" from B to C to E is the basis of melody. The movement here is "virtual" because sounds do not literally move; nothing literally moves *from* one note *to* another, certainly not the notes themselves; but phenomenologically it is *as if* something moves from one position of the pitch continuum to another. This is the "what it's like" of our experience of a melodic line; more than that: there is no "melody" apart from the phenomenal character of such experience. Melody is not

simply “out there”, and neither is it wholly “in our heads”. Like the experience of color, it depends on complex relations of things, of which our cognitive and perceptual biases are one *relatum*. This also suggests that our ability to experience things *aesthetically* is a consequence, among other things, of our *limitations* as cognitive agents. Music depends on our foregrounding of certain sonic features in detriment of others, in virtue of our cognitive and perceptual biases. In this way, being subject to destruction, being *situated* in the world relatively to things that are and things that are not under our power (of imagination, cognition, and action) is what endows things with an aesthetic character.¹⁹¹

Going back to melody and the “virtual movement” of tones; Scruton identifies two essential features of melody “in all its forms”. a) There is an “internal constraint” exerted by every note on every other: “a melody is a sequence in which no note can be altered without changing the character of the whole.”¹⁹² This is a feature not shared by non-melodic sequences of tones; melody gives rise to the “wrong note” phenomenon, our feeling that something is not quite right when one note in a melody is changed. This, of course, can have good or bad aesthetic results. It ranges from a plain mistake to a variation on the original melody. b) Melodies have boundaries, a beginning and an end (“and often half-endings along the way”¹⁹³). Also, the development of a melody can be interspersed, as when its ending is postponed until the end of a section or movement of the work. Melodies are perceived as sonic individuals, contained within boundaries that cannot be identified with any purely acoustic features of the sound event in which the melody is heard. This is a function of our ability to individuate sound sequences within a single complex sound event. From the point of view of the purely acoustic event, there is nothing but a mindless succession of sounds. We perceive melodies because, once again, in our experience of sound some elements are foregrounded and others are made secondary or completely ignored.

Some melodies are more song-like, others more theme-like. Some are more open-ended, like the flowery melismatic vocalizations of plainchant; some form self-contained *tunes* that can “stand on their own”. But these are all variations of one and the same basic phenomenon – a phenomenally distinct kind of acousmatic experience, which consists in hearing a succession of sounds as a “virtual movement” through regions of the pitch continuum; a kind of experience which is an added

¹⁹¹ The philosopher Crispin Sartwell (2004), interestingly enough, characterizes *beauty* as *the object of longing*, and asserts that only beings such as ourselves, i.e., beings who are “in the process of losing everything they have” (including themselves), can experience the world aesthetically. This theory resonates sympathetically, I believe, with Eddy Zemach’s account of aesthetic properties as conceptually connected with desire and aesthetic experience as a *desire-modulated* experience of the look of things.

¹⁹² Scruton, R. “Rhythm, Melody and Harmony”, in Kania & Gracyk (2011: p. 29).

¹⁹³ *Ibid.*

layer to the more “primitive” kind of virtual movement that is rhythm. Scruton also provides a helpful classification of melodies as *logogenic* (melodies that serve the purpose of setting words to music), *orchegenic* (melodies that originate as tunes for dancing and may not be very singable), and *harmonegenic* (melodies that arise from underlying harmonic relations).

Finally, the experience of harmony adds yet another layer to the primeval driving force of rhythm. It is a feature of the acousmatic experience that is most distinctive of “western” music, in that its development depended on the extraordinary intellectual feat of musical notation, allowing music to have a complex organization not only “horizontally”, in terms of sounds following one another, but also “vertically”, in terms of simultaneity. This does not make western music “better” than other musical traditions in absolute terms, although if I may be allowed a degree of political incorrectness, it is certainly a good-making feature (in non-absolute terms), in that it is better to have a tradition that exploits this dimension of musical experience than not having it. It is also better to have musical traditions that explore other dimensions of it in ways that the “western” tradition has not pursued than not to have them, so there is nothing “ethnocentric” about this remark. Besides, “western” tradition has greatly benefited from cultural cross-fertilization – think of how musicians such as Debussy and Stravinsky have not only renovated the “western” tradition by incorporating “foreign” elements, like the delightfully ambiguous whole tone scale¹⁹⁴ and polyrhythms but also taught us, “western” listeners, new ways of appreciating music. And if I may be allowed a personal note on the matter, twentieth-century experiments such as those of Lou Harrison show that what we call “western tonality” is nothing like a “closed” ethnocentric tradition, but a protean beast, with an ever-evolving capacity to find surprising new ways of reinventing itself. Ethnocentrism, unlike beauty (I dare say), is wholly in the eye of the beholder. Musical boundaries are somewhat arbitrary, and constantly being demolished by the inexhaustible joy of musical experimentation. A degree of cultural immersion may be required to properly understand how a certain musical tradition is the way it is, but the joy afforded by the experience of listening with understanding is not bounded by parochial constraints.¹⁹⁵

Just as experiencing rhythm and melody is something over and above the experience of *succession*, the experience of harmony is something over and above the experience of *simultaneity*.

¹⁹⁴ The well-known and already mentioned flute solo from *Prelude à l'Après Midi du Faune* is a good example of this.

¹⁹⁵ Arguing against a view according to which the concept of “musical work” is a late arbitrary imposition of western “bourgeois” culture, Scruton (2009: p. 11) says: “There is a peculiar experience of ‘same again’ which is fundamental to musical organization. There could not be meaningful improvisation without this experience, and the emergence of works from a tradition of spontaneous performance is exactly what we must expect when people listen, and therefore recognize what they hear as ‘the same again’. ‘Play it again’ is the mark of acousmatic joy, and one that is by no means confined to ‘Western art music’ or the concert-hall culture.”

Harmony exists when an aggregate of two or more tones sounding simultaneously is experienced as a single distinct entity. The triads (three note chords) generated on each degree of the scale by adding the second and fourth overtones in the overtone series (counting from the fundamental, which in the case of C gives C, G and E – the C major triad), and the more complex chords one obtains by adding other notes to the lot, are more than mere aggregates. Not all such aggregates form chords. The notes of a chord “go together” in a way that is absent from an arbitrary mesh of tones. In the case of the major triads this feeling of “going together” is reinforced by the fact of the proximity of the constituent notes in the overtone series (a fact on which a poor argument for the primacy of western tonality was once built); but there is also something quite unique in a “dominant seventh chord” or a major triad with an added sixth, or, for that matter, in minor, augmented and diminished triads, and so forth.

Chords “animate” the musical line in a peculiar way, giving it a “momentum”, a sort of “ebb and flow” (patterns of tension and resolution) unlike anything that can be achieved with non-chordal harmony. Some chords are “magnetically drawn” to other chords while others have a more inimical relationship, depending on an immense variety of factors, such as the relative proximity of the keys to which the chords belong (in the “circle of fifths”, which is a graphic representation of how the 12 major and minor keys of diatonic tonality relate to each other in terms of shared notes), but also on stylistic context – e.g., the same chord will sound dissonant in one context (musical *genre* or idiom, individual composer) and consonant in a different context; or it will sound equally consonant in a variety of contexts, though serving different functions, in virtue of which it will sound different in each circumstance (e.g. a C major chord is the tonic in the key of C major, the subdominant in the key of G major and the dominant in the key of F major; this translates into a variety of phenomenally disparate experiences, such as that of the chord in question being the chord to which other chords “strive” to “resolve”, or being the chord that most vividly “compels” movement to the tonic; experiences which are further complicated by all sorts of contextual features). The phenomenon of consonance and dissonance is itself a matter of degree and it is understood in terms of a specific harmonic context.

But harmony does not exist only in music made with *triads* (major, minor, augmented, diminished, sus) plus the chords we obtain by adding further notes to the triad; harmony “is what we hear in two simultaneous pitches when we hear them as a single object.”¹⁹⁶ The oldest forms of harmony in our western tradition consist in the singing of two parallel melodic lines at the distance

¹⁹⁶ Scruton, “Rhythm, Melody and Harmony”, in Kania & Gracyk (2011: p. 31).

of a fourth or fifth (what is known as “organum”) and a melodic line accompanied by a single drone note (usually performed on an instrument such as the hurdy-gurdy). This latter example is the kind of harmony one can find in bagpipe music, which will more closely resemble medieval music since this instrument is tuned to the system of just intonation.

Harmony can also be viewed in two distinct ways, as the driving force that constrains melodic and rhythmic development, in which case we say that a piece has a predominantly “vertical texture”, or it can be the result of the intermingling of two or more independent melodic lines (what is known as “counterpoint”), in which case we say that a piece has a predominantly “horizontal texture”.

The phenomenal characteristic that Scruton isolates as distinctive of the experience of hearing harmony in simultaneous aggregates of tones is that of the said aggregate, whether a dyad, a triad, or a rather complex chord, “filling” the virtual musical space and excluding things from the space it occupies. Chords can be “cluttered” or “hollowed” (as it happens with the “open fifths” or the so called “power chords” in some rock music, which consist in playing the octave and the fifth, eliminating the third).

“Chords are spaced, open, filled, or hollow. They spread over the stave, strain asunder, tend away from or towards their neighbors. They provide the primary experience of a spatial (as opposed to a temporal) Gestalt in music: of a unit which crosses distances and can be grasped all at once.”¹⁹⁷

Scruton’s abundant use of the term “Gestalt” as applied to “musical individuals” is meant to underline the idea that the three “dimensions” along which musical experience develops – rhythm, melody, and harmony – are not reducible “without remainder” to mere relations of succession and simultaneity. Hearing a melody is more than hearing a succession of pitched sounds; hearing rhythm is more than hearing regular patterns of duration and stress; hearing harmony is more than hearing simultaneous pitched sounds. We can accept this idea without committing to the whole of Scruton’s conception. The view that rhythm, melody and harmony are more than acoustic succession and simultaneity is compatible with different ways of conceiving how we perceive these phenomena. Beings with a relevantly different cognitive and perceptual apparatus may hear the successions and simultaneities without experiencing the sounds as tones, i.e., without having what is a musical experience in the primary sense of “musical”. Extended uses of the term “musical”, as when one says about natural sounds (birdsong, a babbling brook) or non-musical artificial sounds

¹⁹⁷ Scruton (1997: p. 45).

(the clicking of wheels) that they are “musical” derive from our ability to hear them *as if* they were made to exhibit “basic musical features” (Andrew Kania’s term for rhythm, pitch, melody and harmony)¹⁹⁸, or, in other words, such uses derive from our ability to *listen for* such basic musical properties in sound events that lack them.

We can either have a theory of music’s nature that includes “organized sounds” that lack the “basic musical features” (e.g. a “symphonic poem” for one hundred metronomes clicking simultaneously, or John Cage’s *4’33’’*) or we can have one that excludes such experiments (this is compatible with accepting their status as artworks of some sort). In any case, I don’t believe that we can have a univocal account of both “central cases” of musical entities (either works or pure improvisations) and the “borderline” cases of *musique concrète*, symphonic poems for metronomes and “aesthetically framed” ambient noises such as *4’33’’*. We must identify a primary sense in which organized sound events (a concept which includes silences as a structural element) are grouped under the concept *music* and an extended, non-primary sense. The phenomenological variety of listening experiences compels it. To experience rhythm, melody and harmony *in* sound events is phenomenologically distinct from *listening for* such basic musical features in sound events that lack them; and it’s also phenomenologically distinct from experiencing a sound event as a dramatic performance or as “soundscape art” (e.g. as an aspect in the design of gardens and fountains) or any kind of artistic exploration of acousmatic experience that doesn’t involve “basic musical features”. Even if we divide the arts of sound in a) poetry and literary performances; b) “soundscape art” /non-musical sound art; and c) music, the latter category can be said to include works that are (or could be) “sonically indistinguishable” from examples of non-musical sound art. I think we should make room for such cases rather than exclude them, and that we should explain their possibility in terms of an unavoidable *procedural* (or institutional) element in the concept of art, as it applies to music. Things that are “musical” in a derived, secondary sense, will have an element of *aboutness* that musical things in the primary sense typically don’t necessarily have (notable exceptions are “musical commentaries” and “quotations” and parodies such as Mozart’s *Musical Joke*). Yet, a simple observation will help us to distinguish clearly between a primary use of the concept *music* and an extended or secondary use. When psychologists and neuroscientists set out to study empirically our responses to music or how music is processed by the brain, they will avoid using “symphonic poems” for metronomes, *musique concrète*, Yoko Ono’s flushing toilet and any sound works that lack “basic musical features”, for obvious reasons. The most exotic specimens

¹⁹⁸ Kania & Gracyk (2011: pp. 7-8).

of non-western musical traditions will be by far less problematic, even if the issue of cross-cultural musical understanding is complicated enough. And this is because we intuitively grasp that musical experience, whether we call it “musical” or not, is phenomenally distinct, and it is *that kind* of experience that is our focus of interest. Our interest in the “borderline” cases, by contrast, is *explained* by our primary interest in this specific kind of experience, quite independently from how we judge them artistically.

Another idea that emerges from all the previous considerations is that an informative account of what music *is* cannot be separated from an understanding of “what it’s like” to experience sounds as music. Any “external” description of music (in terms of the *acoustic* structure of sound, in terms of what “brain resources” we employ in music-making and music-listening, etc.) will only make sense to the extent that we already have an implicit understanding of the points made above concerning acousmatic experience and the “basic musical features”. We begin to understand what music is when we try to clarify what is phenomenally distinct about the experience of music. Hearing music in sounds is akin to experiencing a “visual clash” in a colored surface. We understand what the visual clash *is* only in terms of the qualitative aspects of our experience, since there is no way to produce a semantic equivalence between aesthetic predicates and non-aesthetic predicates (that some overtones are expressed in terms of simple ratios doesn’t *explain* our experience of an harmonic interval as “consonant” rather than “dissonant”; that is, our (non-musical) aesthetic preference for simple ratios doesn’t illuminate our experience of consonance, since, as we’ve seen, harmonic context largely determines what is and is not “consonant”, just like pictorial context largely determines what is and is not “gaudy”, “clashing”, etc.). An exhaustive description of the colored surface plus its “social and historical contextualization” will never substitute the *experience* of the visual clash *in* the colors. We either see it or we don’t. If we don’t then the “external” descriptions will never be illuminating; they only widen our understanding because we know *by acquaintance* the phenomenon we’re trying to understand. This is so in the case of music because *musical properties* (like aesthetic properties in general) are a subclass of *phenomenal properties*. How we conceive of this in terms of basic ontology is, of course, a completely distinct issue.

4.2 The components of musical experience: aesthetic metaphor

The debate over the role of metaphor in *describing* musical experience, and how it connects to the role of metaphor in the aesthetic use of language in general, has important consequences for the definition of music. In particular, it will be determinant to our assessment of concepts such as that of “purely musical experience”, which in turn will provide a verdict on the consistency of “formalist” theories of music (or a certain version of what such a theory consists in). As it happens, I hold that there is no such thing as a “purely musical experience” of sounds, at least not as Hanslick imagined it, since to experience sounds musically *is* to experience sounds in terms of *something else*; and this is no less true of descriptions of sounds in terms of space, motion and action, than it is of descriptions of music in terms of emotion and other expressive predicates (which can be absent without thereby *music* being absent: music can be emotionally neutral). This amounts to a denial of Hanslick’s “positive thesis”, that is, that “absolute” music (music with no words, title or program) is aesthetically superior to non-absolute music (either music with words or instrumental music with suggestive titles or programs). Extra-musical cues such as “Scheherazade’s theme” or “The Old Castle” (from Mussorgsky’s *Pictures at an Exhibition*) in otherwise purely instrumental music are not aesthetically irrelevant or damaging *intrusions* on the music, but one strategy among others of making the aesthetic profile of the music more vivid. Applying “reduced listening” to non-absolute instrumental music, in the sense that one tries to hear the music in abstraction of the extra-musical cues will not result in a more “pure” musical experience, since musical experience is already “impure” at an ontological or foundational level.

It is in the discussion of aesthetic metaphor that apparently independent issues, such as how music relates to the emotions and the issue of music’s definition, turn out to be intimately connected. This also contradicts the more or less standard view that we can make progress in more “specific” issues, such as how music relates to the emotions, while completely ignoring the definitional question. I grant that we can ignore the actual step of articulating a definition, expressed in a declarative sentence of the form “X is Y” (as in fact Scruton does, despite theorizing unabashedly about the nature of sound and the nature of music). What we can’t ignore is the issue of what music *is*, whether or not we play the game of definition and counterexample. As we have already seen, this point is made clear by the very existence of anti-definitional theories of art. When these are relevant at all, they are *metaphysical* theories, that is, they try to explain what *in the nature of art* makes it the case that art can’t be captured in terms of necessary and sufficient conditions. If nothing about the nature of art makes the anti-definitional theory true, then we have no reasons to believe it. Weitz himself, as we’ve seen, argues against the possibility of defining art from a thesis

about an *essential property* of art: that art is such that it is always (there are no instances where it isn't) open to radical change (even in those circumstances where artists are not actually concerned with departing from tradition – e.g. the history of art has registered periods of stasis as well as periods of innovation).

Likewise, whatever the relation between music and the emotions is (whether music *arouses*, *represents*, or *expresses* emotions or whether emotional descriptions of music are always metaphorical) that will have to be explained in terms of whatever *in the nature of music* makes it the case that it can or cannot have such and such relation to emotional states. If the nature of music sets no restrictions on such issues, then any answer to the problem of how music relates to emotions will be arbitrary. We may not need to settle first on necessary *and* sufficient conditions for music in order to determine how music can and cannot relate to emotion, but we cannot avoid getting into the issue of the nature of music, of what music *can* and *cannot* do. To see this one need only note that in inquiring about music's relation to emotions we are inquiring, on the one hand, whether or not (purely instrumental) music *can* arouse, represent, or express emotions and, on the other hand, whether or not it is *essential* to music that it does any of these things, if the answer to the previous question was affirmative. In other words, we cannot help inquiring about music's *modal properties*: what is *always* true of it (like Weitz's open-endedness is always true of "empirical concepts", should he be right in his claim), what is *not always* true of it, and what is *never* true of it. The idea that we can discuss such matters without thereby discussing music's (or art's) nature is one curious example of conceptual illusion.

Going back to "aesthetic metaphor". First of all, what *is* an aesthetic metaphor? An aesthetic metaphor is a metaphorical use of a predicate (any predicate) with an aesthetic purpose. For instance, all aesthetic uses of the predicate "is delicate" are metaphorical. When the word is used literally it means only that the object in question breaks easily if dropped or that it is fragile. Said of a melodic line, it can't be meant literally, since sequences of sounds are not literally fragile. Not all aesthetic predicates are metaphorical; e.g., "beautiful", "ugly", "elegant", "gracious", etc. But the majority of terms we employ with an aesthetic purpose are terms that do "double work" in the language. They are used literally to denote one kind of property and aesthetically to denote either another kind of property or, should one be of a Kantian persuasion (like Scruton), no property at all (they just signal a particular *imaginative* act, like that of seeing aspects in clouds).

Musical aesthetic metaphors are those metaphorical descriptions of music that call our attention to the music's aesthetic profile. This can be understood as referring to a certain way of *imaginatively* perceiving music or as referring to an aesthetic property of the music, one which we experience by experiencing the music's non-aesthetic properties. Philosophers disagree as to what

descriptions are metaphorical. For instance, they disagree over whether *emotional* descriptions of music are metaphorical or literal. If one thinks that the description of a piece of music as “sad” is a literal description, one must then give an account of *how* music can be literally sad: is it because it elicits or *arouses* sadness in listeners? Is it because it *represents* sadness somehow? Is it because it is *expressive* of sadness? If one thinks that the description of a piece of music as “sad” is a metaphorical description, one must then give an account of *what* property, if at all, is being (figuratively) ascribed to music. In a sense, even if we think that aesthetic predicates do not refer aesthetic properties (if we think that an object can’t have aesthetic properties *in the same way* that it has, e.g., a certain width, length and height, a certain surface texture, a certain color, etc.) we are nevertheless ascribing a property to it. For instance, one can think that to predicate delicateness of a musical line is to ascribe to the music the property of being imaginatively perceived in a certain way. This is a genuine relational property, but it differs from other relational properties in some important respects. For instance, a picture of Mount Fuji has the relational property of being *correctly* perceived *as of* Mount Fuji. But it is difficult to say what could make the imaginative perceiving of a musical line as “delicate” a *correct* way of perceiving it imaginatively, if the ascription of the property refers only to the content of the subject’s experience and not to any delicate-making features in the object itself. The cloud really has the relational property of being imaginatively perceived by me as dog-shaped, but there are no *normative* constraints on this kind of imaginative perception. If I happen to see something in the cloud that you can’t see no matter how hard you try (assuming I’m being sincere), what in the cloud could make my imaginative perception *inappropriate*? And would I have the right to say *your* experience of the cloud is somehow *defective*? But this is not the way we tend to behave when it comes to *aesthetic* descriptions. If I see as heavy and awkward something you see as gracious and elegant, you won’t be able to repress the idea that something must be wrong with my eyesight, even if you otherwise pay lip service to aesthetic subjectivism. So, a normative aspiration is part and parcel of our experience of things has having an aesthetic profile; the question is whether those aspirations are grounded or not. Do we both suffer from an illusion when I see as heavy and awkward the *same* object you see as gracious and elegant? Are we unwittingly translating social biases into a ludicrous kind of fictional language (aesthetic language)? Are you insidiously trying to “oppress” me by making me conform to a standard of “taste” which is nothing but the epiphenomenal manifestation of how we both stand in relations of power quite alien to the look of things and how it relates to our perceptual and cognitive abilities? This is not the place to delve into that discussion, but, needless to say, I think that the kind of “sociological skepticism” about aesthetics represented in the latter questions is dead wrong, *whatever* turns out to be the right view about aesthetic normativity.

In the last section I tried to explain a view that sees music as the object of a phenomenally peculiar kind of experience of sound events. According to that view, music is not just “organized sound”, but a peculiar kind of organization that we perceive in sounds, when we perceive sounds *as* sounds, as opposed to perceiving sounds as (nothing but) the acoustic indicators of their causes. The description of that experience was given in terms of spatial and motion metaphors: rhythm “pushes forward” and gives the sequences of sounds in which it is heard a peculiar “demeanor” (it makes music move “with a physiognomy akin to ours” – which is likely why philosophers in antiquity saw music as exerting either a morally uplifting or erosive influence); melodic lines move “up” and “down” and “through” an imaginary space, not identifiable with the space in which the sound event is produced (no melodies “travel” through the space occupied by performers and their instruments; what literally happens in that space is a continuous disturbance of the air; an amorphous mesh of vibrating patterns); harmonies inflect the rhythmic and melodic movement and “fill” the space through which the melodic line travels, giving us a sense of “boundary” and “progression”, of “departure” followed (or not) by a “return” to familiar shores. In Scruton’s words, “we hear music as spatially organized, and if we did not do so we would be unable to understand the art of music as we know it.”¹⁹⁹

Scruton’s thesis is controversial, because it can be understood in more than one way, as making a more or less demanding claim. No one disagrees that music *can* be described in terms of spatial and motion metaphors; but there is disagreement on whether such descriptions are *essential*, in the sense that it would not be available to us the phenomenally distinct kind of experience referred to as “musical experience” if we did not *hear* music *as* spatially organized in the ways described. Malcolm Budd (2003), for instance, has challenged the idea that metaphors are essential in this way to the description of music. In his view, descriptions of music in terms of movement are *literal* descriptions, to the extent that they denote a movement which is *purely temporal*, and this would be the kind of movement that characterizes the experience of sounds as music. But it is not at all clear whether the notion of a “purely temporal” movement is not itself metaphorical. If it is, what it is a metaphor of is just *succession*. We can illuminate this idea with a simple thought experiment. In certain fantasy movies we see characters endowed with the superhuman power of “stopping time”. What this actually means is that the character in question makes everything around him “freeze” while he is still able to move freely. Not only time keeps “passing” (another spatial metaphor?) but there is physical change in such scenario: the character moves around and interacts with the

¹⁹⁹ Scruton, R. “Rhythm, Melody and Harmony”, in Kania & Gracyk (eds.) (2011: p. 31).

environment in several ways. Now imagine a scenario where the whole universe is “frozen”, with no superhumanly endowed character messing around while everything else remains still. About that scenario, either we are allowed to say, in Budd’s terms, that there is also a “purely temporal” movement – a succession of “moments” in which nothing changes, nothing moves through space – or we have to say that there is no such thing, since temporal succession is conceptually linked with physical change. If nothing changes, time does not “pass”. In that case, “moving temporally” is just a metaphor for change. Accordingly, when we say something like “the melody moves from B to C to E” (in a “purely temporal” way) that means we experience the successive stages of a sound event, or, in other words, the “unfolding” of the melody. But as we have seen from the contrast between acousmatic and non-acousmatic listening, *all* sound events “unfold” (i.e. as events they have temporal parts) but not all sound events afford the experience of melodic movement. So experiencing a melody moving “up” and “down” the pitch continuum is merely to experience the unfolding of a sound event whose temporal stages consist of pitched sounds (B, *then* C, *then* E). But, as we have seen, it is possible to hear such successions *without* hearing any musical “movement” – such musicless unfoldings of sound are most likely what our pets experience coming out of our stereos, not to mention the case of tone-deaf people (some amusiaks are unable to identify melodies without sung words and others report experiencing something not unlike the chaos of pots and pans colliding in what for us is, say, a musically coherent D minor symphony).

This is how I would adjudicate the Scruton-Budd controversy: Budd is right in claiming that there is a literal sense in which music “moves”. This is the same sense in which we say that events “unfold” (though “unfold” is being used metaphorically here – and a curious metaphor at that, since all *literal* unfoldings are events), that is, the temporal stages of events “follow” one another in succession, and since all sound sequences are events, all sound sequences “move” in this literal sense. But this is not the sense in which we normally say that the music “moves”. When we speak of musical movement we do not speak of the temporal succession of sounds, but of something we hear *in* that succession of sounds. It is more akin to the sense in which we say that “the story is rising to a climax” or “heading to a closure”, or the sense in which we say that someone is “going down to the point”. It is something that occurs *through* the unfolding of the sound event or *by way of it*; it is not the unfolding itself. And in *this* sense, our use of spatial and motion predicates as applied to music is a metaphorical use. Take the case of the succession: B *then* C *then* E. This succession is equally compatible with

a) describing the melody going “up” from B to C to E (assuming the interval between B and C to be a minor second, and the interval between C and E to be a major third);

- b) describing the melody as going “down” from B to C to E;
- c) describing the melody as going “left” from B to C to E;
- d) describing the melody as going “right” from B to C to E;
- e) describing the melody as going “north” from B to C to E;

... and a plentitude of other imaginable description in spatial terms.

This is because there is no *literal* sense in which a sound of “higher” pitch than another (e.g. one with a frequency of 659 cycles per second, relative to one with a frequency of 523 cycles per second and yet another with a frequency of 494 cycles per second)²⁰⁰ is in the *spatial* relation of “higher than” with another sound, except perhaps in the uninteresting sense that the sound source of one sound can be placed spatially above the sound source of another sound. But we don’t hear the music moving “down” if we hear a note produced on a bass clarinet followed by a note on the piccolo flute, even if the former instrument is positioned above the latter. Neither has “upward melodic movement” to do with the increasing number of cycles per second, since “number of cycles per second” does not figure in the qualitative character of our experience of an upward melodic movement. We don’t *experience* sounds as patterns of cycles per second (the same way we don’t experience colors as wavelengths), even though that is what physically underlies our experience of the sounds – if one presses the central A key in a tuned piano, one will hear a sound *produced* by something (a string) vibrating at 440 cycles per second; if one then presses the A key an octave above, one will hear a sound produced by a string vibrating at 880 cycles per second. But the phenomenally distinct experience of “higher” that ensues is not *about* the number of cycles in a vibratory pattern.

So, my conclusion is that while there is a literal sense in which music “moves” in a purely temporal manner (in the sense that all instances of music are events and events have temporal parts that *follow* one another), that fact is perfectly compatible with the truth of the following two propositions: a) that there are metaphorical descriptions of music in terms of space and motion and b) that such metaphorical descriptions are ineliminable in the sense that without them we have no resources to describe the experience of music. The idea is that the same motion predicate can be used both literally and metaphorically as applied to music, and in each case the predicate expresses a different *concept*, and both concepts pick out different *properties* of the music. The concept

²⁰⁰ These frequencies correspond to B, C and E, according to current standard tuning.

expressed by the literal use of motion predicates picks out the property that stretches of music share with all events as things that have temporal parts; the concept expressed by the metaphorical use of motion predicates picks out the property in virtue of which we hear sounds as music, i.e., as having rhythmic, melodic or harmonic implications.

In this way, we could grant Budd's contention that all metaphorical descriptions of music are *in principle* permutable with literal descriptions without thereby acknowledging that metaphor is dispensable. And this can be understood in two ways. a) *Permutation of use* – the metaphorical predicates can be used literally in such a way that they don't refer to the same property when used metaphorically. Permutation of use is not "truth-preserving" since the sentences in which the predicate is differently used actually express different propositions. b) *Permutation of predicates* – we can replace the metaphorical predicates with different words used literally as long as both predicates pick out the same property. Predicate permutation is truth-preserving if the same property is picked out in both cases (if the concept expressed is the same). Truth-preserving predicate permutation would only serve Budd's point *if* our ability to use the non-metaphorical predicate was *independent* of our use of the metaphorical predicate to pick out the property in question; for instance, if we could become proficient in picking out the property we describe as "musical sadness" (the sadness we hear in the music) while *never* thinking of that property in terms of sadness. Imagine using "glurg" to refer what we usually mean when describing music as "sad". "Glurg" would be intelligible enough while thinking about it as a surrogate term for (musically) "sad"; but is it conceivable that we would pick out the *same* musical feature with "glurg" if we had *never* come to think of music in terms of sadness? It just doesn't seem possible to represent or think of the property we usually refer to as "musical sadness" without thinking of it in terms of sadness.

Since there are theories that view the relation between music and emotions such as sadness in literal terms – theories that seek to make plausible the claim that music is somehow *literally* sad – metaphor would be dispensable at least in the case of *some* emotional descriptions, should those theories turn out to be true (should that be the case, a further question, of course, would be whether *all* expressive descriptions of music could be accounted that way). But in the case of "musical movement" (and other descriptions of music in terms of space and motion) the opponent of metaphoricism would have to show either that there is a literal use of motion predicates that picks out the same property picked out by "musical movement" or that the property picked out by "musical movement" could be independently picked out by different predicates (other than

space/motion predicates), in their literal use, such that we could pick out that property with those predicates without ever having thought of the relevant properties in terms of space and motion.²⁰¹

Partly, Budd's resistance to the thesis that descriptions of music in terms of space and motion are *essentially* metaphorical is due to his resistance to the controversial thesis, held by Scruton, that metaphor is not limited to our *descriptions* of music but somehow enter the *perception* of music. This is known in the literature as "the metaphorical perception thesis" and a version of it is defended by Christopher Peacocke in the more limited domain of music's *expressiveness* – which includes but is not limited to *emotional* expressiveness. Peacocke (2009) relies on the notion of a *cross-domain mapping* or isomorphism that is exploited in some perceptual states, which he calls "experiencing something *metaphorically as* something else". To elucidate this, he lists some kinds of *experiencing something as*:

- 1) Experiencing F as F – this occurs, e.g., when one sees a cat as a cat.
- 2) Experiencing F as G – this occurs in perceptual illusions, such as when one sees a cat under a dappling light as a raccoon.
- 3) Experiencing R as a representation of F.
- 4) Experiencing F metaphorically as G.

The first two kinds of "experiencing as" correspond to all cases of perception of something under a *sortal*. Some of those experiences are veridical and some are illusory, depending on whether the conditions of observation have us applying the right or wrong sortal.

The third kind of experiencing as corresponds to cases such as when one sees a marked surface as a *depiction* of a cat or when one experiences a non-figurative set of marks, e.g., a word token, as a (non-depictive) *representation* of a cat.

The fourth kind is not a case of perceptual illusion nor, according to Peacocke, a case of *imagining something to be something else*, as when I imagine the inanimate objects laid out on the table in front of me as if they were animated objects (people, animals, etc.). Most objections raised against Peacocke claim that he is really talking of imaginative perception of the latter kind, and that there is no such perceptual state as *seeing metaphorically as*. Peacocke counters such remarks on the basis of phenomenological vividness. If we are to be faithful to the phenomenology of such perceptual states, we must distinguish them from other cases where imagination plays a role in

²⁰¹ For an insightful discussion of these issues, see Nick Zangwill (2011), especially §2.

perception but that lack the phenomenological vividness he refers to. Two examples he presents are a painting by Francisco de Zurbarán (*Still Life With Pottery*), a depiction of kitchen utensils, in which some pots are endowed with distinct anthropomorphic qualities, inducing in the viewer an experience of the depicted objects as if these were people; and another painting by Caspar David Friedrich (*The Solitary Tree*), in which an isolated tree is depicted in such a manner that we experience it as a solitary person. Though both are examples of figurative representation in paintings, Peacocke argues that the kind of perceptual experience involved is not essentially of pictures. We can have the same kind of experience with non-depicted, real life pots and trees, or any other object. A common example of this is the experience of willow trees as *sad*, which is a property no willow tree can literally exemplify, since trees have no mental states. According to Peacocke, such perceptual states “exploit” a cross-domain mapping or isomorphism between the domain of objects to which the actually perceived object belongs, and the domain of objects capable of literally instantiating the properties which we see the perceived object as having. In the willow tree example, we experience the willow tree as instantiating properties that no tree can literally instantiate, thus making it the case that the willow tree is experienced *metaphorically* as a sad human being.

Peacock applies this idea to musical expressiveness in the following way: when we listen to a piece of music such as Debussy’s *Voiles* (Sails/Veils) we are actually able to hear something “fluttering” in the music, that is, we experience the musical movement *metaphorically* as a fluttering movement, the movement of fluttering objects, such as sails or veils. It may be that in some cases the listener must be given some extra-musical cue (such as a suggestive title) in order to perceive this in the music, but even where extra-musical cues play a role in “guiding” the experience of music, there is a marked phenomenal difference between those cases where the cues are effective (e.g. when we do in fact hear the fluttering *in* the sounds) and those in which they are never effective, no matter how hard the listener tries to experience the music in accordance with the extra-musical cue. There are also cases where we need no cues at all, as when we perceive sadness or exultation or some non-emotional expressive trait (such as *fluttering*, *submersion*, etc.), the same way we don’t usually need a linguistic cue to experience a willow tree as sad.

One standard objection to all of this is that metaphors are a *linguistic device*, so that to say that a perceptual state is itself metaphorical is nothing short of a category mistake. The perceptual states that Peacocke is interested in, an objector might hold, are rather appropriately described as cases where imagination is more or less *compelling* or unavoidable rather than subject to the will, unlike what happens with linguistic metaphor. Peacocke counters by pointing out that a state of perceiving metaphorically as can be elicited by an exercise of the will (as in the case where I initially don’t see

the Zurbaran pots as people but in the course of trying to see them in that way, exercising my imagination, something “clicks” and, all of a sudden, I *am* in fact able to see them in the intended way). Such “exercises of the will” are not successful most of the times (cases like the Zurbaran pots or Friedrich’s solitary tree are uncommon). This objection strikes me as weak, since it is open for us to say that trying to see the object according to the cue (which can be a metaphorical *description* of the object as such and such) sometimes results in compellingly imaginative perception. Peacocke argues that linguistic metaphors are themselves dependent on our ability to have perceptual states that exploit a mapping or isomorphism between disparate domains of objects. However, it is hard to shake off the idea that we could equally and as easily say that our ability to have perceptual states that exploit such mappings is dependent on our ability to make metaphorical descriptions of objects. Maybe that is a proper subject for empirical enquiry rather than *a priori* reasoning. Nevertheless, both ideas strike me as false. I’m inclined to believe that it is more plausible to think that both the ability to have such perceptual experiences and the ability to make linguistic metaphors are independent feats of imagination that equally depend on a common ability to exploit the so called mappings between disparate domains of objects. Take the well-known negative metaphor “No man is an island, entire of itself”. On a Davidsonian view of metaphor²⁰² (with which I am sympathetic), the sentence has no “metaphorical meaning”, only the literal meaning given by the (literal) meanings of the words used (if “island” was used with a meaning other than its literal meaning, the metaphor would not be the same metaphor it is). The metaphor makes us notice something, a property of human beings that is quite distinct from the property of not being a stretch of land surrounded by sea (as it happens, the negative metaphor is also literally true). How could my ability to understand this metaphor, and notice what it calls my attention to, depend on my ability to perceive people “metaphorically as” *not being* islands? It does not seem to depend even on my ability to see the victim of a shipwreck lost in the ocean “metaphorically as” an island, supposing this would be possible with the same phenomenal vividness that Peacocke appeals to in his other examples. I am able to understand the metaphor even if I *never* have imaginative perceptual states involving a transfer of concepts between the domain of people and the domain of insular entities. So it seems that the most plausible view on this is that we have an underlying ability – that of exploiting (perceptually or otherwise) mappings or isomorphisms between domains of objects – and this ability manifests itself in cases of compelling imaginative perception *and* in linguistic devices such as metaphor, but these are independent feats of our cognitive apparatus. As Stephen Davies as

²⁰² See D. Davidson (1978).

remarked about this issue: “The mechanism Peacocke proposes – cross-domain transfer and isomorphism – could underlie and make apt my imagining. If so, there is no need to invoke a special metaphoric mode of perceiving or, for that matter, a special metaphoric mode of imagining.”²⁰³ This is precisely what I mean by the same ability or “mechanism” underlying both imaginative perception and figurative uses of language, while the latter remain distinct, independent phenomena.

Whether or not one finds Peacocke’s ideas on “perceiving F metaphorically as G” plausible, they make explicit what is implied in saying that metaphor “enters” the perceptual experiences in which sounds are experienced as musical sounds, and is not restricted to the metaphorical *description* of such experiences with words. And even if we don’t accept that there is such a thing as a metaphorical perceptual state, our brief journey through Peacocke’s theory allows us to conceive in a more plausible way what Scruton is aiming at with his idea that metaphor “enters” the perception of music “at the foundational level”²⁰⁴ (that is, that metaphor permeates the perception of sounds *as music*). It is important to note that Peacocke never proposed his theory of “perceiving metaphorically as” as an account of how sounds are perceived *as music*, but merely as an account of music’s expressive qualities, that is, as an account of how sound events which already qualify as music are also perceived under certain concepts, namely, those expressed, e.g., with emotion predicates (though not all expressive predicates are emotional predicates) used metaphorically (predicates for properties that music cannot *literally* exemplify).²⁰⁵ Unlike Scruton’s, this is *not* intended as an account of musical perception “at the foundational level”.

However, Peacocke’s view can be used to give more clarity to Scruton’s proposal on how sound events become musical sound events, even if we end up rejecting the notion that there are metaphorical experiences or metaphorical perceptual states. To see this, it suffices to note that to hear “virtual movement” *in* a sound event is always to hear a certain *kind* of movement. There is no such thing as an aural experience of movement *per se*, as if we were able to experience an abstraction of movement. If we hear “virtual movement” in music, that is always a peculiar kind of movement, e.g., a fluttering movement, as in Debussy’s piece.

It is not only friends of metaphoricism who appeal to the mechanism of cross-domain mapping or “conceptual transfer”. Particularly, a defender of the resemblance theory of musical expression,

²⁰³ S. Davies (2011: pp. 22-23).

²⁰⁴ The use of the phrase “at the foundational level” to refer the perception of sounds as music occurs in Budd (2003: p. 222).

²⁰⁵ As I’ve explained, this is compatible with there being legitimate *literal* uses of the same predicates that express different concepts.

James Young (2012: p. 392) has claimed that “we experience music as moving as a result of cross-domain mapping.” Young (1999; 2001) had previously argued for the view that musical expressiveness is best understood in terms of *representation*: music is expressive of sadness and other emotions because music represents those emotions. According to Young, music is able to represent emotions in virtue of the intentional exploitation, by the composer, of our ability to hear movement in the sounds, to hear music as moving. In his theory, resemblance operates at two levels. At the “raw” level, the experience of music resembles the experience of motion. By intentionally exploiting this, the composer makes it the case that music represents motion, for instance, Debussy’s piece represents the fluttering motion of sails or veils. This is a case of *direct representation*: movement is the direct object of the representational intentions of the composer who intentionally exploits the resemblances between experience of music and experience of movement. Once we are able to experience these resemblances, we are able to experience further resemblances between the experience of different kinds of movement and the experience of behavior expressive of emotion, namely, *bodily movement* and *vocal expression*. In this view, then, musical representation of emotions is *indirect*: music directly represents movement and indirectly represents emotion, by way of the two-stage resemblance mechanism between experience and experience of movement, on the one hand, and of expressive bodily and vocal movement, on the other.

There are two especially relevant points to take notice here: a) regardless of any authorial intentions to exploit the resemblance mechanisms in order to represent movement, action and emotion, the resemblances may be there to be experienced in the music (e.g. we can experience the music as “fluttering” even if there was no intention to directly represent fluttering movement and indirectly represent something else); b) whether or not we endorse the view that descriptions of music in terms of movement, action and emotion are metaphorical, whether or not we endorse the view that there are such things as metaphorical perceptual states or “metaphorical experiences”, or whether we endorse a resemblance view of musical expressiveness (which is typically averse to metaphoricism)²⁰⁶ the same psychological mechanism of “cross-domain mapping” is appealed to by defenders of all these views.

This means that even if we reject the Scruton-Peacocke idea of a kind of perceptual state which is metaphorical in nature, in favor of an explanation of musical expressiveness in terms of resemblance, this is compatible with Peacocke’s suggestion that the relevant perceptual experiences

²⁰⁶ Young, no less than Budd, considers descriptions of music in terms of movement and emotion as literal, non-metaphorical.

(whether of the anthropomorphic pots in Zurbaran's painting or of expressive qualities in music) involve an element of "conceptual transfer" or "isomorphism". The same psychological or conceptual mechanism can be invoked by an explanation of musical expressiveness in terms of resemblance. (Another way to put this point is that the same perceptual states described by Peacocke as "metaphorically hearing as" can be redescribed in terms of experienced resemblances, without any change to the phenomenal character of such experiences.) The point that a resemblance theory is not only compatible with the (essentially?) metaphorical character of emotional descriptions of music but may even be used to explain them was more recently made by Schroeder (2013).

A plausible sense in which metaphor "enters" the perception of music "at the foundational level" (explaining what makes an aural experience a *musical* experience and thus what a sound event must have to be a *musical* sound event) is not one in which we do not perceive sounds as music unless we produce certain *descriptions* of them in metaphorical terms (we can hear music without attempting to describe it), nor is it one in which the perceptual states in question are somehow "metaphorical" themselves or are "guided" by a linguistic metaphor (the metaphorical descriptions are elicited by the experience, not the other way around); it is rather the sense in which if we set out to describe what makes an aural experience a *musical* one, we will unavoidably make use of metaphorical descriptions. And this is because even if we could substitute literal predicates for the metaphorical ones, we would not be able to think about the properties in question except in terms of the concepts expressed by the metaphorical predicates. In other words, even if we coined new (literal) predicates to describe "basic musical features", rather than speaking of chords as "filling" a virtual or imaginary space, of melodies as "traveling" through "regions" of the pitch continuum, etc., those predicates would not be intelligible as descriptions of the music apart from our ability to think about the sounds in terms of space and motion.

Significantly, Saam Trivedi has argued against Scruton in the following terms:

"*Contra* Scruton, the distinctiveness of a timbre might be heard under very different metaphorical descriptions or under none at all; for example, the literally shrill timbre of an oboe holding a high note might be heard as such even by little children incapable of understanding metaphors. And one can hear melodic or musical movement without appealing to Scruton's spatial metaphors. For example, a melody can be heard as moving from the leading note to the tonic in the familiar musicological terms of melodic tension and resolution (or melodic drive or yearning) that we

literally hear in the music, or in some such terms that describe the experience without essentially referring to spatial features.”²⁰⁷

Scruton's thesis would be vulnerable to these objections if understood as claiming that the relevant perceptual states are dependent on the actual articulation of the description of music in linguistic terms. But little children capable of hearing the shrillness of the oboe's timbre are in the same epistemic position of an adult who can hear the shrillness while ignoring the meaning of the word "shrill". We search for words in order to describe the experience, we are not enabled to experience the features described by the predicates as a consequence of learning the words (it is possible to be in a somber mood without being able to use the word "somber" to describe that mood). If the opposite were true, it would seemingly make mysterious how we are able to learn to apply words correctly. When we pick up the word "shrill", learning how to use it correctly in future instances, what happens is that we converge or coordinate with others in the qualities involved in the experience of shrillness and become more proficient in describing it. As our vocabulary grows, we become able to describe our experiences more accurately (even if non-literally), whereas before we would only be able to do it but in a more awkward or fuzzy way. Otherwise it would be a complete mystery how we can, upon learning a new predicate, have the distinct impression that the term is "just right", as opposed to some other, to serve our purposes. Phenomena such as wide social convergence in the use of aesthetic predicates like "shrill", "somber" or "mournful" would be nothing short of a miracle. Some people describe the oboe's timbre as "nostalgic" or, to use our previous example, the *duduk*'s timbre as "mournful". We don't become capable of hearing such qualities in the sound *after* we learn the word, just as we don't acquire the ability to experience sadness only after we have learned how to use the word "sad", or to experience the qualitative aspect of pain that is described by the metaphor "a stabbing pain" only after we learn the word "stabbing", nor, for that matter, are we able to perceive objects as red only after we learn the predicate "is red". Of course, the widening of one's verbal resources, our ability to describe the world in increasingly more fine-grained terms, may inflect the way we experience the world, particularly, it can make our musical experiences more rich, complex and discriminating, which in its turn can widen the boundaries of our experience of the world in general. Our experience of symphonies, string quartets, concerti and sonatas (all variants of the same musical form) is not the same after we learn the concept of "sonata-form". Even if we are not able to subsume the

²⁰⁷ Trivedi, S. "Music and Imagination", in Kania & Gracyk (eds.) (2011: p. 117).

experience of a stretch of music under that concept, because it takes time and effort to follow large-scale form instead of focusing exclusively on a moment-to-moment experience of the unfolding sound event, we will (in many cases at least) try to focus not only on the more ephemeral connections but on how the musical movement coalesces and coheres as a whole. Someone who ignores the concept of sonata-form will of course be able to experience a melody appearing later in the movement as a restatement of something heard before, but that experience will have a different quality after we learn the concept. This parallels the way our attention to the features of a painting will be differently focused after we acquire concepts of stylistic elements and historical and contextual factors. It is not that we can't experience the "emotional tone" of a painting or some other expressive feature before we have the linguistic tools to describe it more aptly, but the more aptly we are able to describe it, the more discriminating our experience in turn becomes.

As for the description of the music in the "familiar musicological terms", I believe I have already convincingly argued against the idea that such descriptions can be independently meaningful in the absence of phenomenal qualities that the musically trained listener shares with the "musical illiterate". If we observe the mechanism of a clock while ignoring what a clock does, the purpose with which it was built, the "cogs and wheels" description of what is happening will leave us in the dark. A description in terms of "moving from the leading note to the tonic" only makes sense to anyone with minimal musical training because he is acquainted with (has first-person knowledge of) what the expression means. Furthermore, the "technical" description has no point other than to translate the distinct experience of movement from the most "unstable" note in the scale to the "gravitational center" of the melody built on that scale. And we need such technical descriptions because the metaphorical description in non-technical terms is not useful as a "blueprint" to replicate the effect in different contexts. But that it is of no use for that purpose doesn't mean it is not apt at describing *what is going on* when one "moves from the leading note to the tonic". If the phenomenal qualities were not there, the technical description would be akin to the description of a mechanism that serves no purpose, like cogs and wheels connecting in such a way that no purposeful movement results. What gives a leading note its distinctive character is the fact that it "leads" to the tonic, and this phrase stands to the phenomenal character of the melodic movement in the same relation that "shrill" stands to the perceived quality of the oboe's timbre that makes the use of *that* word apt. We *never* would perceive Scheherazade's Theme in Rimsky-Korsakov as *martial* in character in virtue of hearing it described that way, no matter how often.

In fact, most technical vocabulary used to describe music is riddled with metaphor. Take "scale" for instance, which is a spatial metaphor used to describe something that literally has no "steps" nor does it "ascend" or "descend" from one point to another. Or think of the already given example of

plagal cadences (a major or minor triad on the subdominant preceding the tonic); the expression “plagal cadence” perhaps contains a hidden metaphor, since “plagal” comes from an ancient Greek word for “oblique”. There is already a suggestion of movement with the term “cadence” (to fall) and the term “plagal” is added to describe the inflected character of that movement, which makes sense as a musical description in virtue of how its referent is phenomenally distinct from other cadences. Of course, one should beware of etymological considerations; we must recognize that contemporary uses of a given word may be only faintly connected with the original meaning. For instance, “plagal” might be interpreted literally, according to its ancient Greek meaning, in which case it would contribute to the spatial metaphor (as applied to sounds). More plausibly, though, it just followed the terminology already in use to contrast the “authentic” modes of plainchant with the “plagal” ones (always beginning a fourth below the *finalis* of its authentic counterpart).²⁰⁸ Most people who use the term just use it (regardless of any etymological considerations) to refer to that cadence which “bypasses” the dominant chord with its characteristic leading note or, what is the same, that is not an “authentic” perfect cadence (as the cadence from the dominant chord to the tonic is known) and so does not afford so strong a sense of closure. Anyway, the point is that it must trace back to some phenomenological aspect of the music. There is nothing especially “musical” about the phrase “plagal cadence” that would be absent from any other (figurative) description of the musical event in metaphorical terms, given by an attentive but “musically illiterate” listener. Technical vocabulary is more fine-grained, and more suited to serve the “blueprint” function, but it is as “impure” as any “poetic” description of the music. Yet it is precisely this residue that makes it intelligible musically. A completely “depurated” vocabulary would perhaps be suited to the description in purely acoustic terms of the sound events where musical movement figures, but it would leave out the music precisely because the music is what shows itself on the surface of linguistic description by way of that figurative residue.

The other descriptions Trivedi provides as “literally heard” in the music are so obviously metaphorical that the point hardly requires counterargument: “tension and resolution” and “melodic drive” or “yearning” are in fact *literally heard* in the sense that “heard” is being used literally. But the qualities we literally hear are being described metaphorically as “tension”, “resolution”, “drive” and “yearning”: it is *we*, with our cognitive and perceptual biases, that hear a dominant seventh chord as “yearning” for the tonic and being “resolved” in it. Sounds themselves do not “yearn” nor do they “resolve” into anything, though this is how their “aesthetic surface” is made manifest to us.

²⁰⁸ There were eight “church modes”, consisting of four authentic/plagal pairs.

Granted, Trivedi is merely arguing for the conclusion that *spatial* metaphors are not ineliminable from our basic experience of sounds as music. His idea is that the qualities thus described don't have to be thought of in terms of spatial movement. Budd offers a similar argument to the effect that the qualities Scruton describes as "movement in a virtual or imaginary space" can be equally described with the metaphor of "systole and diastole"²⁰⁹ (the experience of rhythm) or in terms of something "contracting and dilating", but it seems to me clear that our ability to think of rhythm in such terms is dependent on our ability to think of a rhythmic sequence in terms of space and motion. After all, "contracting and dilating" are events that can only occur where something occupies different regions of space at different times and the metaphor of "contracting and dilating" is not likely to be apt for *all* cases of rhythmic movement. That we can point to the same features of music that compel description in terms of space and motion by using different metaphors doesn't show that we *don't* have to think of those features in terms of space and motion if those other metaphors are to be intelligible; and that the literal use of "contracting and dilating" applies to things that move in space points in the opposite direction. My guess is that we won't understand descriptions in terms of "purely temporal" movement (succession) if not for our previous understanding of events as things that exist in space and not merely in time. It makes sense to speak of "movement" from a temporal stage of an event to the following stage only because the event itself consists in things that are gaining and losing properties in more than a "purely temporal" sense. We can imagine, e.g., the gradual change of color in a leaf as a temporal "unfolding" only because we imagine patches of color occupying different spatial parts of the leaf at different times. If nothing is going on "in space", we might as well say that "time has stopped".

Let us focus once more on the phrase "movement from the leading note to the tonic". This could be taken as a non-spatial, "purely temporal" movement: the leading note is followed by the tonic. Has we have seen, even if Budd is right in there being a literal use of "movement" that refers to a mere temporal succession, this is too wide a net to capture the phenomenology of melodic movement, since we can hear the same succession of leading note and tonic in different contexts, without thereby hearing "the same thing". Think of the experience of hearing those same notes at the beginning and at the end of a melody. A description merely in terms of succession is not enough to capture what is going on and neither does the description "leading note-tonic" makes it any clearer, for in each token succession of the *same* note types we experience a *different* movement. To further complicate things, these two successions of the same elements will have a different

²⁰⁹ Budd (2003: p. 221).

character depending on the melodic, harmonic and rhythmic context, and my guess is that if one tries to describe the phenomenally distinct character of the experience in each case, one will be compelled to make it more clear by using either a spatial/kinetic metaphor or a metaphor involving the concept of things occupying space, moving or traveling through space, filling a space, and so on. And, in any case, substituting metaphors for metaphors will not help Budd's or Trivedi's case against Scruton, since it is not enough to deny that spatial/kinetic metaphors are indispensable; they must deny that any metaphorical description is essential. And, as we've seen, successfully paraphrasing metaphorical descriptions in literal terms will not of itself eliminate whatever in music makes its description in terms of space, motion and action metaphors phenomenally compelling, since the literal surrogates are only intelligible *as surrogates* for the phenomenally compelling descriptions. In particular, "technical" descriptions of music show themselves, on close inspection, to be intelligible only in virtue of the former. There would be nothing "musical" about a taxonomy of chords or cadences, if not for a distinctive kind of aural experience, with its peculiar phenomenological character.

However, I agree with Trivedi in that we don't need to conceive of metaphor as somehow "built-in" our perceptual states to account for the nature of musical experience. There is nothing metaphorical about the perceptual states themselves. While it is true that such perceptual states "exploit" mappings between domains, the same is true of the experience of seeing a maple tree or Mount Fuji in a picture, without there being a need to invoke a special kind of "metaphorical" perceptual state to explain the experience of figurative depiction. The fact that the perception of music is imbued by concepts not literally instantiated by sound events gives us no reason to postulate a special kind of perceptual state. If by "metaphorical perception" one means only that the perceptual state "exploits" a mapping between domains, then one is no longer talking of metaphor in the sense that we say that "No man is an island" is a metaphor. One is then applying the term "metaphor" to a more general phenomenon, one which also underlies the making of metaphorical descriptions but is not subsumed in the later concept. Cross-domain mapping is involved in producing metaphorical descriptions and in the phenomenon of imaginative perception, but all are conceptually distinct things. That our perception of a marked surface is "informed" by the concepts of a dog or of Mount Fuji does not make the perceptual experience a "metaphorical" one. The same is true of our musical experiences and of the anthropomorphic pots in Zurbaran's painting.

A more significant factor in Budd's rejection of the essential character of metaphorical descriptions of music has to do with an "objection of principle". The objection of principle is that

“music cannot have a merely metaphorical essence”.²¹⁰ The truth of this claim would render incoherent Scruton’s thesis that basic musical perceptions “involve” spatial metaphors as their constituents. The argument for the claim, that makes explicit the objection of principle, runs along the following lines:

“Even if the arguments Scruton advances were sufficient to establish an essential relation between the experience of melody and the concept of movement . . . to call the description of the nature of these experiences by means of the related concepts metaphorical is not to specify the alleged relation: the relation between the experience and the property signified by the concept, or the way in which the concept informs the experience of the sounds, is not illuminated by the notion of the metaphorical transference of the concepts or properties to the experience. For any metaphor stands in need of interpretation; a metaphor is susceptible of indefinitely many interpretations; and an interpretation of a metaphor is the injection of significance into it, not the extraction of significance from it.”²¹¹

The argument is not immediately transparent, so I’ll try to clarify it as I see it, whether or not this exactly matches what Budd explicitly says in the relevant passages. If I’m correct, the dispute between Budd and Scruton will turn out to hinge on a more general problem concerning how the meaning of metaphorical sentences is to be understood.

There is a gap between what a metaphorical sentence *means* and what we mean by making a particular metaphorical utterance. We use words metaphorically when we apply a predicate (a word) to an object that cannot literally exhibit the property referred to by the predicate. A consequence of this is that any affirmative metaphorical sentence will typically be literally false²¹² while negative metaphorical sentences (such as “No man is an island”) will be literally true, since the latter states that the predicate does not apply to the object, which will render the sentence literally true in virtue of what is typically the case with metaphor. Yet, what we (or the poet John Donne) mean by the metaphor “No man is an island” is not to point out the fact that no man is a stretch of land surrounded by water but rather something like “No human being is a self-sufficient

²¹⁰ In Tanner& Budd (1985: p. 242).

²¹¹ *Ibid.* p. 241.

²¹² Nick Zangwill (2014) gives the example of an affirmative metaphorical sentence which is also literally true. Imagine that someone has two jobs, one as a teacher and another one as an electrician. Now imagine that someone says of that person that “he brings light to the world”. The sentence will be literally true if taken as a reference to his work as an electrician, but also metaphorically apt if taken as a reference to his work as a teacher. Cases such as this account for the qualification “typically” above. But the vast majority of affirmative metaphorical sentences will be literally false, such as “Juliet is the sun”.

entity”. But any literal paraphrase we produce to explain what we mean with a particular metaphor is not what *the sentence* means. The meaning of the sentence is given by the meaning of its constituent words only. It follows that an indefinite number of literal paraphrases (as well as an indefinite number of non-literal paraphrases) will be equally apt as an interpretation of the metaphor; but then no specific paraphrase can be given as *the* metaphor’s non-literal meaning. The meaning of the metaphor just is its literal meaning, and what we mean by making a metaphorical utterance will be explained in terms of pragmatic force. By uttering a metaphor, one intends to call attention to or point out some real feature of the object described metaphorically. We point to a cross-domain resemblance or mapping: human beings are not “disconnected” from one another, in the way that islands are disconnected from one another and from a mainland by intervening masses of water. We understand a metaphor when we grasp what his utterer is pointing at. But what the utterer is pointing at is independent of the metaphor, that is, the metaphor is not *constitutive* of the features it is used to call attention to.

This sheds some light on Budd’s objections to Scruton’s thesis. If metaphors are not constitutive of what they direct our attention to, then in principle any feature pointed to with a metaphorical predicate can be referred to with a literal one. In principle we could introduce literal predicates to replace all predicates used metaphorically. But if this was the heart of the objection, it would be too easy. The problem runs deeper and it has to do not just with the features themselves that we point at using metaphors but with our epistemic limitations in picking out those features. In the case of “No man is an island” we could bypass the metaphor and use one of the available literal paraphrases to “make the same point”, even if that would not make John Donne’s poem aesthetically any better. In many cases, the use of metaphor is not compelled by any epistemic limitation in the way we are able to think of the features referred to, but by the fact that language is put to an aesthetic use, over and above its use in conveying propositional thought. This doesn’t mean that both uses must conflict cognitively. The aesthetic exploration of cross-domain mappings, e.g. as it occurs in poetry, can make us notice things that are not so conspicuous when language is not used to perform such “double work”.

But if literal meaning is *primitive* then this seems to give force to Budd’s objection that music cannot have a “merely metaphorical essence”. This would be tantamount to reversing the order of explanation. The concept of metaphor cannot be used to illuminate the concept of musical experience, since the metaphors with which we characterize musical experience will themselves be in need of clarification, something that seemingly can only be achieved through literal paraphrase. What this means is that unless there is a primitive, literal sense in which music is said to “move”, the spatial metaphors will be as mysterious as the concept of musical experience they are meant to

illuminate. This is what Budd means (I think) when he says that the appeal to metaphor can't specify the relation between the experience and the concepts deployed in the description. That specification would have to be achieved by undoing the metaphor via paraphrase; and so Budd's verdict is that Scruton's position is incoherent.

An interesting point in all this is that Budd acknowledges the possibility that an essential relation between the experience of music and the concept of movement might be established. The problem lies not in the idea that we experience music in terms of movement but in the idea that metaphors are "embedded" in the experience of music, that is, in the idea that a metaphor or a metaphorical thought somehow figures in the perceptual states that make up the experience of music. This prompts the question: why doesn't Scruton restrict his metaphor thesis to the linguistic description of music? Why the extra step of making metaphor part of the experience? My guess is that this problem is rooted in Scruton's anti-realist approach, not only to basic musical features but to aesthetic qualities in general. In Scruton's view, aesthetic metaphors, of which the descriptions of music in terms of space and motion are a subclass, do not refer to properties in a robust sense of the term. Metaphorical descriptions, on this view, do not attribute properties to the music. As such, they do not have genuine truth conditions. Descriptions of music in terms of space and motion, as well as descriptions of music in terms of emotion and other aesthetic/expressive predicates, do not pick out mind-independent properties. Music is not "out there" in the world but exists only in our experience of sounds, within the "theatre of the mind". To hear the music as sad is not to perceive a property that sounds exemplify, since sounds are mindless, insentient things; neither do sounds literally move. To hear the music as sad or a melody as moving, jumping, plunging, sliding, falling, soaring, etc. is akin to see a cloud as animal-shaped. Aesthetic metaphors express non-propositional imaginative acts in which things are perceived under aspects. Basic musical features, in this view, are not features of the sound events in which they are heard, but features of our *experience* of the sounds. This is why Scruton is compelled to make metaphor a part of our experience of music, instead of restricting it to our *talk* about music. Doing this would prompt the question: what properties are being referred to then? Scruton cannot specify those properties since he thinks that *no* properties are being genuinely ascribed. There are no aesthetic facts serving as the truthmakers for aesthetic descriptions of music. Aesthetic language does not aim at truth.

“it seems to me that in our most basic apprehension of music there lies a complex system of metaphor, which is the true description of no material fact. And the metaphor cannot be eliminated from the description of music, because it is integral to the intentional object of musical experience.”²¹³

There is an argument in Scruton for the ineliminability of metaphor from the experience of music and not merely from the description of music, that takes the form of a dilemma: take the example of musical sadness; a description of music in terms of sadness cannot be attributing to the music the same property we usually denote with the term “sadness”, since sadness can only be possessed by conscious organisms. But neither is it *another* property, since words used in metaphors retain their literal meaning. There is no “metaphorical meaning”, only metaphorical *uses* of literal meanings (Scruton’s conception of metaphor actually predates Davidson’s). And what would be the point of using the word “sadness” if it was being used with an unrelated meaning? From this Scruton concludes both that metaphor is ineliminable *and* that no property is being ascribed to music. Musical sadness is not in the music, it is wholly a part of the phenomenology of our experience. Hearing music is not a matter of perceiving mind-independent properties but of the “active creation of an intentional world, in which inert sounds are transfigured into movements, harmonies, rhythms – metaphorical gestures in a metaphorical space”.²¹⁴ The dilemma presented above is intended to apply beyond the realm of music, as part of an argument for aesthetic anti-realism more generally. And this is why Scruton is prevented from countering Budd’s objections with an interpretation of the metaphors: because the theoretical boundaries of his anti-realism force him to see the metaphorical thoughts supposedly involved in the perception of music as a kind of bedrock. Of course, it is an open question whether there is an adequate anti-realist response to Budd that eschews Scruton’s thesis of the ineliminability of metaphor, and I want to leave that door open.

Stephen Davies (2011)²¹⁵ has given an ingenious critique of the thesis of the essential metaphorical character of the description of music. It casts a reasonable doubt over the special role ascribed to metaphor by other theorists and is also a sobering reminder that there is nothing magical about metaphor. Davies’ approach gives central stage to polysemy as opposed to metaphor. Take once again the case of musical sadness. The problem with explaining what property is being attributed to the music by use of the predicate “is sad” is that sadness seems to be integral to the phenomenology of music. We have a vivid experience of the music as sad, as if the sadness of the

²¹³ Scruton (1983: p. 85).

²¹⁴ Scruton (1983: p. 100).

²¹⁵ Especially chapter 2, “Music and Metaphor”, pp. 21-33.

music was not relational, just as we experience colors as of the objects that have them, as if possession of the color by the object was not actually a relational fact. Any explanation of musical sadness must be faithful to the phenomenology of the experience (which does not mean, of course, that we must think of musical sadness as a non-relational fact). The explanation must account for the vividness of the experience. This is precisely the difficulty faced by theories seeking to explain the emotional expressiveness of music in terms of make-believe, by which we imagine the music as the expression of our own emotional states or as the expression of some fictional *persona* that we imagine as “inhabiting” the music. The obvious difficulty is that we don’t need to actually entertain such thoughts in order to experience the music as sad. The same difficulty is faced by expression and arousal theories: music’s expressiveness is not merely a symptom of the composer’s emotional states, since the composer need not be in a sad mood in order to make sad music. And neither can we easily identify musical sadness with the *arousal* of sadness in listeners by music, since the music doesn’t have to be sad in order to elicit sadness. I can be made sad by hearing all kinds of expressive music that no one would describe as “sad”. The sadness of sad music seems *independent* of both the emotional mental states of the composer and the listener. This is not to deny that the acts of making and listening to music can’t be *expressions* of one’s emotional states. I can express my sadness by playing sad music; though I can also express my sadness by playing joyful music. There seems to be no way of going around the problem of phenomenological vividness concerning the independence of music’s expressive properties.

Consider also the view according to which music is sad because it somehow *represents* sadness. This is relevantly different from an expression theory, since representations of sadness don’t need to be taken as symptoms of experienced sadness in order to represent their subject. Neither do they require that someone who recognizes what is represented in a representation of sadness to be sad in order for the representation to be a representation of sadness. Representational properties may be mind-dependent in the sense that nothing is a representation unless someone *intends* it to be a representation *and* it can be recognized as a representation. As Goodman has compellingly argued, resemblance is neither necessary nor sufficient for representation, so that for X to be a representation of Y it is not required that X resembles Y. Resemblance is a symmetrical relation, that is, if X resembles Y then Y resembles X, but that X is a representation of Y doesn’t make it the case that Y is a representation of X. Resemblance is also reflexive. Everything resembles itself, but that doesn’t make it the case that everything is a representation of itself. Resemblance is also cheap: between any two things there will be many resemblances in some aspect; but that doesn’t make it the case that any given object is a representation of any other object it resembles in some aspect.

But this, it seems to me, also provide us with a reason not to identify musical sadness with some representational property of the music. This is not to deny that music *can* represent sadness. Rather, the point is that even if music can represent extra-musical reality, musical sadness does not consist in the fact that music represents sadness. One reason for this is that even if music can be said to *resemble* sadness in some aspect that will be insufficient for it to be a *representation* of sadness. Furthermore, if music can have representational properties, it seems it could (at least in principle) be used to represent sadness without resembling it. In fact, similar cases of musical representation are cheap. Take the familiar example of Scheherazade's Theme from Rimsky-Korsakov's "tone poem". The theme *represents* Scheherazade *because* it is used for that purpose and because suitably informed listeners will recognize it as such. There is a common objection to music's ability to have representational properties in the same way that the word "Scheherazade" represents Scheherazade by denotation. If Rimsky-Korsakov or someone else didn't provide us with extra-musical cues to guide our appreciation of the music, the theme would not represent Scheherazade. But this is surely a bad argument. In virtue of which does the sound/squiggle "Scheherazade" represent its target? Surely not in virtue of some intrinsic representational property, since *all* representational properties are *relational*. Words have representational properties because they constitute stable points of convergence in behavior; they are solutions to a coordination problem: that of representing certain items in the world. Words do not possess a magical power to represent whatever they represent. If people can converge in this way in the use of the sound/squiggle "Scheherazade", why couldn't they likewise converge in the case of Scheherazade's musical theme? That I would be unable to know that the theme refers to Scheherazade if I ignored the music's "program" does not show the theme lacks representational properties, any more than my ignorance of Japanese makes it the case that the expression "wabi-sabi"²¹⁶ has no representational properties. Or take the case of Tchaikovsky's *1812 Overture*, which is meant to represent the battle of Borodin. The composer makes use of national anthems, namely, Rouget de Lisle's *Marseillaise* to denote the Napoleonic army, and a Russian folk tune to denote the Russian army. We cannot know this just by hearing the music, without some knowledge of historical and contextual facts, but this is not relevantly different from the fact that I can't know what "wabi-sabi" means without knowledge of some historical and contextual facts.

At the same time, we can see why representational theories of musical expressiveness won't be satisfactory. The simple reason is that representation and phenomenology can diverge, even if

²¹⁶ "Wabi-sabi" is an aesthetic term, denoting a specific kind of beauty. See Crispin Sartwell (2004, chapter 5).

sometimes the phenomenal aspects of the vehicle of representation are conscripted to make the vehicle perform its representational function (e.g. music which resembles sadness can be used as a representation of sadness). In other words, sad music can be used to represent sadness, but representation is not the mechanism, or at least it is not the whole mechanism, by which we *experience* music as sad. Representational properties are an addition, an extra layer in the cake, so to speak. A picture that represents Napoleon does not do so merely by resembling him in certain aspects. We don't have to accept Goodman's view that all representation is conventional in order to see that the ability to recognize a picture as of Napoleon also depends crucially on knowledge of historical and contextual facts: that there was the individual named "Napoleon"; that the picture resembles him; that the maker of the picture intended to represent him, and so on. This doesn't require us to do archaeological work, of course. In most cases, the relevant contextual properties (including the maker's intentions) can be inferred from the experience of the picture, given that we already have the required background historical knowledge. But pictures, like words, have no inherent magical power to represent their targets, apart from how we *use* them to perform their representational functions.

Does this favor in any way the position according to which descriptions of music in emotion and other expressive terms must be metaphorical? In his case against Scruton, Budd argues that even if certain descriptions of music start as metaphorical, this is due to the fact that we initially lack literal terms to describe the features of the music we want to describe. But with time, the metaphors die out; they become dead metaphors (as was the case with many other expressions in use, such as the term "legs" applied to a chair) and are used as literal extensions of the primary meaning. This seems to be the case with many descriptions of timbre. Typically, timbres are referred to via their sources ("an oboe sound"), but this soon proves insufficient, since there can be considerable timbral differences between note tokens of the same note type, played on two instruments of the same kind (e.g. two oboes), or even between tokens of the same note played differently on the same instrument. Consider the already given example of the "mournful" tone of the *duduk*. Even granting that "mournful" is a more illuminating description of the phenomenal character of that instrument's timbre than a description of it by the name of the instrument, it seems true that we don't have to *think* in terms of mournfulness to (re)identify that phenomenal character; we are able to grasp it independently of that particular metaphorical description, and neither do we mean by it that the sounds produced with that instrument have emotional states. In fact, that "mournful" aptly describes the instrument's tone color doesn't make it the case that the instrument *always* sounds in a way that fits the description, no more than an oboe's sound is always shrill, poignant and nostalgic or a clarinet's sound always bright and playful. What happens is that in the absence of a literal term to

specify the phenomenal character, we resort to figurative language, and this initially figurative use becomes a literal extension of the primary use, as in the familiar case of dead metaphors. In this way, in speaking of “musical sadness” the case would not be relevantly different from having had a literal term to pick out that phenomenal aspect of music from the start. We could have used “glurgness” as a literal term for musical sadness, instead of “sadness”. In this view, the current musical use of the term “sadness” is no more *essentially* metaphorical than the use of “legs” is when applied to parts of tables and chairs. A dead metaphor is dead precisely because the term is now used in the same way that a literal term would be used if adopted from the start to the same purpose. The dead metaphor of musical “sadness” is, in this way, a surrogate to our counterfactual “glurgness” that *could* have been adopted to pick out the same phenomenal aspect of music. This is a view according to which the concept of musical “sadness” is homonymous but not synonymous with the usual concept of “sadness”.

The problem with this strategy is that musical sadness doesn’t seem at all to be identifiable in a way that is disconnected with the normal use of the concept “sadness”. The convergence in descriptions of music in terms of “emotion, motion and height” is too pervasive for this to be a mere quirk of language. The absence of literal terms to pick out the phenomenal aspect in question does not of itself explain convergence in the adoption of a particular metaphor as an alternative. The fact that the same feature can be pointed at by using many disparate metaphors suggests that the use of any particular metaphor has an element of arbitrariness that doesn’t seem to adjust to the phenomenology of musical experience in the case of predicates for emotion, motion and height. These are too pervasive and compelling for it to be the case that the musical feature we refer to by “musical sadness” is not crucially related to our previous grasp of non-musical, garden variety sadness. It is just not the case that an emotional metaphor was introduced to supplement the lack of a literal term, where any other metaphor would be equally apt, as long as we are able to (re)identify the phenomenal feature in question. It seems rather that the phenomenal feature in question is not graspable at all independently of our previous grasp of the common meaning of “sadness”. The same applies to predicates of space and motion.

It is at this point that Davies’ strategy seems more plausible. He points to the fact that, unlike “live” metaphors, which require creativity from the part of those who coin the metaphor and have an element of arbitrariness, in that a different metaphor could be used to point out the same feature, “musical sadness” displays the important characteristic of its meaning being preserved across languages. Even if “sad music” does not display the same exact features in every musical culture, there is nothing culturally awkward in the cross-cultural use of a sadness predicate to pick out that quality in music. This, according to Davies (2011: p. 26), seems to make the phenomenon of

emotional descriptions of music more akin to polysemy, which is typically preserved in translation, than to live metaphor, which is typically not so preserved. Likewise, descriptions of music in terms of space and motion are shared cross-culturally, so that the use of such predicates cannot be arbitrary in the same way that using this or that particular melodic theme to represent an extra-musical object has an element of arbitrariness to it. In the case of musical sadness, there is a secondary use of the predicate “sad”, Davies argues, that is deeply related to its primary use, as when one says that basset hounds have “sad faces” or that the mask of tragedy has a “sad look” or as when we say that someone “looks sad” without thereby implying that he actually feels sad. It is in this way, according to Davies, that we describe music as “sad”, that is, in virtue of perceived resemblances. There is a legitimate extension of the primary sense of the term that applies not to psychological states themselves but to physical appearances or “contours”. The same is true of predicates of space and motion as applied to music: by describing a melody as “slow”, “stumbling”, “drooping” and the like we do not imply that music literally moves in space, but that the phenomenology of musical experience makes those descriptions compelling, or, in other words, the experience of music resembles the experience of movement, even if the music does not actually move through space. The experience of musical movement exhibits relevant *resemblances* to the experience of spatial movement without it being the case that music literally moves in space, just as the “sad” appearance of willow trees does not imply that willow trees feel sadness. These are neither homonymous concepts with distinct meanings nor are they live metaphors, but literal extensions of the primary meanings of predicates for emotion and spatial movement.

“When I coin a novel metaphor, I creatively project something I imagine onto the world and invite others to do the same in working out what I mean. But a person’s description of music as sad does not involve a novel, creative use of language that provokes others to try to puzzle out what she could be driving at, given how incongruous her choice of words appears to be.”²¹⁷

At this point, I want to step back and look at what is at issue. Why does it matter whether descriptions of music in terms of emotion, motion and height are literal extensions of the primary use of such predicates rather than (particularly compelling) metaphorical descriptions of music? Nick Zangwill, who endorses the thesis of the essential metaphorical character of descriptions of music in terms of emotion, motion and height, disagrees with Scruton that the predicates ascribe no properties at all to music and he disagrees with Scruton’s idea that there is something metaphorical

²¹⁷ S. Davies (2011: p. 32).

about the *experience* of music itself. He refuses the idea that descriptions of music in emotion terms are literal (even in an extended sense) because, following in the steps of Hanslick, he refuses that a literal relation between music and emotions accounts for music's expressiveness. Sad music is not sad because it arouses, expresses or represents emotion. He puts music's (emotional) expressive properties on an equal footing with all its other *aesthetic* properties. In this way, descriptions of music as "sad", "joyful" or "gloomy" are not relevantly different from descriptions of music as "ethereal", "delicate", "vibrant", "powerful", "shimmering" and a host of other aesthetic predicates. He draws here on Frank Sibley's pioneering work on the nature of aesthetic concepts. In a well-known paper on aesthetic concepts, Sibley notes that the vast majority of "aesthetic" predicates have a primary non-aesthetic use, and that the case of "beauty", "ugly", "elegant", and the like, which have a primary aesthetic use and are figurative when their use is non-aesthetic, are the exception, not the rule:

"Clearly, when we employ words as aesthetic terms we are often making and using metaphors, pressing into service words which do not primarily function in this manner."²¹⁸

Thus, the term "delicate" in its primary use refers to a non-aesthetic property of objects: their fragility or tendency to break when dropped. But when "delicate" is being used to describe the appearance of a Ming vase, of a statue or of a melody, the property denoted cannot be the same that the predicate denotes in its primary use, for if we had a way of making a Ming vase unbreakable without altering its appearance, this would not make it look any less delicate, and sounds cannot break nor be dropped, so the "delicateness" of melodies cannot even be understood in terms of an unbreakable object actually looking like something that *would* otherwise break easily if dropped, since sounds cannot be dropped nor break.

Zangwill's (2014; 2015) solution to this is to claim that predicates used aesthetically do not express the same *concepts* as they do in their primary, non-aesthetic use. So we have one predicate – "delicate" – and two different uses of that predicate: a primary, non-aesthetic use and a metaphorical, aesthetic use. In each use, a different concept is expressed by the predicate, namely, delicacy_{NA} (non-aesthetic delicacy) in the first case, and delicacy_A (aesthetic delicacy) in the latter. (Zangwill explicitly says that such conceptual difference is not to be seen as a difference in the meanings of words. The meaning of "delicate" does not change when the use of the word shifts between aesthetic and non-aesthetic.) At this point, Zangwill owes us a story of how the two

²¹⁸ Sibley (1959: p. 422).

concepts are related, so as to avoid Budd's objection that the same feature could have been independently picked out by a literal predicate, distinct from *delicacy_{NA}*. He opts for a causal story, in which our ability to use the predicate to pick out the aesthetic property is causally dependent on our ability to grasp the non-aesthetic property (much like our ability to walk is causally dependent on our previous ability to crawl). In this way, we would not be able to grasp the music's delicacy, or its sadness and other expressive properties, if we did not have a previous grasp of the non-aesthetic concepts expressed by those predicates in their literal use.²¹⁹ But what accounts for this causal connection? What compels us to describe music's *delicacy_A* and music's *sadness_A* with the same predicates we use to pick out *delicacy_{NA}* and *sadness_{NA}*?

Here we reach, Zangwill argues, the limit of what can be said with the philosopher's tools. There is an agreement here with Peacocke's view on musical expressiveness, in that the conceptual mappings or isomorphisms appealed to by Peacocke in explaining how music is experienced as sad, joyful, somber, gloomy, etc., cannot be further elucidated by *a priori* thought. The philosophical explanation ends with the appeal to that mechanism. The question of what mappings are "psychologically real" or vivid is an *empirical* question, not a conceptual one. The philosophical explanation is that in hearing a "fluttering" movement in Debussy's *Voiles* we have a perceptual experience that exploits a mapping between those things to which the non-aesthetic concept "fluttering" applies and the domain of musical sounds. *Why* that experience is phenomenologically vivid, or why some mappings are more appropriate to our experience of the music than others, is no longer within the reach of philosophy. Davies seems to point in a similar direction.

"When something is simply our shared way of experiencing music and sound, rather than talking of imagination and projection, it is more to the point to consider how what we all perceive is structured and filtered by our evolved perceptual, cognitive, behavioral, and affective systems."²²⁰

I can't help but doubt whether anything crucial is at stake in the dispute over the metaphorical nature of music's description in terms of emotion, motion and height. Both in the view that such descriptions are metaphorical and in the view that they are literal extensions of the predicate's primary use there seems to be a relevant conceptual difference – musical sadness and musical motion are not the same as non-aesthetic sadness and motion – but not conceptual disconnectedness. In both views there must be a link between the distinct uses of the same

²¹⁹ See Zangwill (2015: ch. 4).

²²⁰ S. Davies (2011: p. 32).

predicate. Both views seem to rely on the same mechanism: our experiences of music exploit mappings between disparate domains of objects, and this happens in virtue of our cognitive and perceptual biases, as well as in virtue of the music's objective features, on which our shared experience of the music converges. It is this convergence that allows us to acquire and correctly apply aesthetic predicates to the music. What then is gained or lost by opting for "metaphoricism" over "literal expansion" or the other way round? The crux of the matter, it seems to me, is in preserving a distinct *aesthetic* use of the predicates for emotion, motion and height. Here is how Zangwill briefly formulates metaphoricism about musical description.

"Although we describe music using words for emotion, motion or height, this is not because real emotions, motions or heights are there in the music, or because the music stands in some relation to real emotion, motion or height."²²¹

But this seems also to apply in the "literal expansion" view; here we also have conceptual difference while preserving word meaning. After all, those predicates which have a primary aesthetic use (beautiful, elegant, etc.) are not metaphorically applied to music, but they still refer to aesthetic qualities of music; and neither would this change if we came to know that some of our literal aesthetic predicates in fact derive from dead metaphors. To the extent that the "literal expansionist" seems to hold that predicates for emotion, motion and height, as applied to music, are akin to dead metaphors, the crucial point, which is the unity of a theory that groups descriptions of music in terms of emotion, space and movement together with other aesthetic descriptions of music, seems unaffected by our choice between both approaches. It all seems to bottom out in a classificatory detail, of whether most aesthetic descriptions are to be thought of as "live" or "dead" metaphors. Both views are closer to each other than to Scruton's in a more important aspect, which is that of aesthetic descriptions actually referring to qualities of the music and not merely to features of our experience of music, as it is the case with Scruton's anti-realism.

If I'm correct, what did we gain in discussing the relative merits of metaphoricism and literal expansionism? At least this point was made clear: that if the motivation for metaphoricism is to avoid what seems to be a consequence of "literalism", namely, that properties such as musical sadness are to be thought of in terms of a relation between music and "real" states of sadness, a relation specified in terms of *expression*, *arousal* or *representation*, then I think this result is

²²¹ Zangwill (2015: p. 80).

secured in both views. As long as the phenomenal aspects of music to which we refer with predicates like “sad” are seen as independent from the music’s power to express, arouse or represent emotions the defender of an aesthetic definition of music will have no problem in granting that music *can* express, arouse or even represent emotions without thereby accepting that to do these things is *essential* to music. Rather, the expression, arousal and representation are additions or extra layers in the cake. These will be things that music can do without it being the case that it is music *in virtue of* doing them, nor will musical experience be *musical* in virtue of being an experience of expression, arousal or representation. The phenomenal qualities of music in virtue of which it *is* music can be responsible, in some cases at least, for the arousal of emotional states; or they can be conscripted to perform expressive or representational functions, in much the same way that one can use a sad-looking mask to express his feelings of sadness. To experience the music as sad, or as “moving through an imaginary space”, will be independent and explanatorily prior to those further experiences of music as expressive, arousing or representational.

4.3 Music as the art of tonal sound

Our exploration of different theoretical approaches to the experience of music and its description in terms of space, motion, action, emotion and other qualities allowed us to devise a way in which a “natural-kind” theory of music can be elaborated, as opposed to “cultural-kind” theories of the procedural or institutional sort. This is the idea that the nature of music is best elucidated by a core kind of *experience*, which is the experience of sounds as *tones*. Sounds are experienced as tones when they are perceived as having rhythmic, melodic and harmonic implications, as well as timbral features that are relevant to its structure and expressive profile. We have seen that the precise nature of the experience of sounds as tones can be specified in different ways, analogously to how the experience of *seeing-in*, that some consider to be at the heart of figurative depiction, can be specified in different ways. The crucial point is that there is a difference to be noted and that the nature of music is not to be understood primarily in procedural or institutional terms, no more than the nature of depiction is to be understood in terms of the same kind of conventions that explain semantic representation. Just like the central experience of *seeing-in* separates the waters without forcing a particular view of how we are able to “see things” in marked surfaces (which is an empirical question), the central experience of hearing sounds as tones separates the waters between a natural-kind view of music (or, if one prefers, causal-agentive functionalism) and a cultural-kind view of music (or, if one prefers, proceduralism or status-functionalism). Hearing sounds as tones, that is, experiencing a sound event as organized to exemplify rhythmic, melodic, harmonic and timbral features (or, for short, “basic musical features”), is the NKA on which an aesthetic theory of music relies. This is an extension to the particular case of music of a more general aesthetic-functionalist theory about art, according to which something is art if its causal history includes an aesthetic insight that certain aesthetic properties will be sustained by certain non-aesthetic properties and an intention to make an object that has the relevant non-aesthetic properties. The view, however, doesn’t presuppose aesthetic functionalism about art. One could in principle be an institutionalist about art while being an aesthetic functionalist about music. Aesthetic functionalism about music is the view that something is music if it was organized to be experienced as having “basic musical features” (where the apparent circularity is eliminated by explaining “basic musical features” in terms of rhythmic, melodic, harmonic and timbral properties) and its relevant non-aesthetic sound properties sustain the envisaged aesthetic (musical) properties. In this view, “basic musical features” are thought of as parameters for the production of determinate aesthetic properties. The account of music I’ve been delineating is thus diametrically opposed to a conventionalist view such as Goodman’s (which can be classified as a “cultural-kind” theory in

Dickie's terms). Music is *not* a system of symbols nor does it belong to such a system. Music can be *conscripted* to perform symbolic functions (e.g. when there is convergence in treating a theme in a symphonic work as standing for a fictional character) but these imposed functions are not constitutive of music nor of musical experience.

I want to revisit a remark by Stephen Davies that social and artifactual kinds such as music are relevantly different from natural kinds such as that referred to with the term "ruby". Rubies are what they are (a variety of aluminum oxide) independently of human beings believing it is the frozen blood of the gods, he says. But music is different. Music has no "universal and unchanging essence".

I believe this statement encapsulates a deeply ingrained conceptual illusion. The first component of this illusion is the idea that the essential properties of something are, just because they are essential, the most important or relevant of its properties *for us* (and conversely, that those which are most important or relevant for us are its essential properties). The second component is the thought that since what is most important for us in music is flexible or variable across cultures and times, then music's essential properties must differ across cultures and times. Now, while it is true that what is most important for us in music is flexible in this way, it is false that essential properties are the most relevant to us. The properties in virtue of which a sound event is a *musical* sound event probably never figure in a list of what is most relevant to us in music, the same way the essential properties of metals and stones do not figure in a list of properties in virtue of which those things are important to us. If one bothered to make such a list, one would only find in it surface properties, phenomenal properties, and such. While it is true that we worry about, say, whether this sample of yellowish malleable metal is really gold, or whether this red stone is really a ruby and not a spinel, the properties in virtue of which humans came to value such things and ascribe them all sorts of culturally flexible "meanings" are surface properties, precisely, those properties that figure in the dependence base of those objects' aesthetic properties. But this is no different from our concerns about whether a painting is an original and not a forgery, or whether the beautiful sounds we hear were really organized by a human agent. A perfect replica of a painting will arguably share all of its aesthetic qualities, unless of course, our knowledge that painting *P* is a forgery somehow alters its aesthetic properties. Opinions are divided here and intuitions are insufficient to settle the matter. But this is not the issue that I want to focus on. The point is that those things that are important to us in music and that are flexible across times and cultures (and even within the same "culture") are *surface* properties of the music, properties that *depend* on its *essential* properties. Medieval harmony is much different from the harmony of eighteenth and nineteenth century "classical" music and both are quite different from the harmony of jazz (things that would sound awfully dissonant to

a medieval listener and perhaps to a modern listener of medieval music are perfectly consonant and expected in jazz) but all depend on the same basic kind of experience of sound: simultaneous tones are perceived as “filling” an imaginary space and endowing the music with a peculiar kind of movement or momentum.

But this conceptual illusion of the “changing essence” is a quite compelling one. It is manifested in the purely lexical solution that some theorists have found to deal with music’s kaleidoscopic variety, namely, in the tradition of talking about world *musics*, instead of a single cross-cultural activity. But here we stumble against the obvious problem: what makes the several “musics” *musical*? What makes a musical tradition a *musical* tradition, as opposed to a non-musical tradition of sound art? To deploy the plural concept *musics* is, precisely, to deploy the concept *music* in the plural. It is like speaking of *homines sapientes* instead of *homo sapiens*, on account of the wide phenotypical differences between human beings. Just like in the case of fallacies (arguments that *seem* either valid or sound but in fact are not), which become more or less obvious depending on the particular argument, the conceptual illusion I’m describing is a kind of reasoning that seems more obviously wrong when we apply it to some class of objects instead of others. In the case of biological taxonomy, it has none of the appeal that makes it so common when applied to social and artifactual kinds.

We have seen how proceduralists about art think that “indiscernible duplicates” reveal the true nature of art by putting “intrinsic properties” aside and making us focus on relational properties. We have also seen that the contrast between intrinsic and relational is a confused one, since, on the one hand, relational properties can be intrinsic (my DNA profile is intrinsic to me, but I have that property in virtue of other things being the case) and, on the other hand, it is hard to think about a genuine case of a non-relational property. The idea here is that all properties involve relational elements, even when we refer to them using monadic predicates. The right contrast is not between the intrinsic and the relational, but between the intrinsic and the extrinsic, on the one hand, and the relational and the non-relational, on the other.

The proceduralist about art reasons in the following way: if two distinct artworks can be perceptually indistinguishable, and if an artwork can also be perceptually indistinguishable from a non-artwork, then the defining properties of art can’t be among the object’s perceptual properties. So far so good, but the mistake is then to suppose this shows that a proceduralist theory is automatically rendered the best candidate. The proceduralist will argue that the relevant properties must be status-functional (we arrive at this clarification of the proceduralist theory by examining the nature of institutional reality); but causal-agentive functions are also non-perceptual. Two perceptually indistinguishable objects will share the same aesthetic properties, at least those

aesthetic properties that depend solely on perceptual ones. But there is a non-perceptual difference between two perceptually indistinguishable objects, where one is a non-artifact that has certain aesthetic properties and the other is an artifact that was made with the intention that it should have those aesthetic properties. Aesthetic *functions* are as relational and non-perceptual as status-functional properties of objects. So the “indiscernible duplicates” methodology does not give a *prima facie* advantage to a procedural theory of art. I have argued previously that the functionalist/proceduralist contrast is unclear, in the sense that *both* families of theories appeal to functional properties, but of different kinds. We can now see how the distinction is potentially confused in another sense: the imposing of causal-agentive functions is no less *procedural* than the imposing of status-functions. The difference is not that for one kind of theory the object must go through a certain procedure while for the other kind of theory it doesn’t; but that the procedures are of different kinds. Both causal-agentive artifacts and status-functional artifacts have their functions in virtue of procedures; the difference is that causal-agentive functions are crucially dependent on the object’s dispositional properties or causal powers while status-functions crucially depend on convergences in representation among social agents.

Nevertheless, the argument from “indiscernible duplicates” is a valuable heuristic tool. We can use it to draw a distinction between hearing something *as* music and something’s being music. Imagine a naturally produced sonic event that is perceptually indistinguishable from a non-controversial piece of music. Both sonic events will be different to the extent that the latter is an *artifact*, which is characterized by certain non-perceptual properties: it has the *function* of exemplifying aesthetic properties that depend on sounds. The former will share some or many of these properties but it is not music, though it can be *perceived as* music and, while perceived in this manner, will appear to share *all* of the latter’s aesthetic properties.

Granted, western music is different from non-western musical traditions. There are also many different musical traditions within the so-called “western” tradition. But it is much more plausible that African drum music, the Arab *maqams*, the Persian tradition of *radif*, the Balinese and Javanese gamelan music, and native-American chants as well as European and North-American “classical” instrumental music, are different manifestations of the *same* basic activity, than to think that these are essentially different things that, inexplicably or irrationally, we feel compelled to describe with our concept *music*. It is intriguing that the same kind of anti-essentialist ideas doesn’t arise in the case of depiction. There doesn’t seem to be a disagreement about whether aboriginal rock art or Paleolithic cave painting contain *depictions*, even if those cultures had no explicit concept for *depiction*. What depiction *is* is independent of whatever concepts, if any, are used by people who make depictions by producing marks on surfaces. It is curious that terms such as “art” and “music”

elicit that kind of anti-essentialist confusions, whereas terms like “depiction” or “painting” do not. I venture that the reason for this difference is that “art” and “music” bring with them all sorts of honorific or evaluative connotations, whereas “depiction” and “painting” do not, and that fact generates a host of conceptual illusions, such as the ones I’ve been focusing on. Concerns about the absence of a “stable essence” (one wonders what an *unstable essence* is: is it exemplified by those base metals that would turn into gold in the alchemist’s crucible?) can easily be seen to be confused if we think about the case of ancient Greek *music*. There is no doubt that the ancient Greeks *had* music (in our sense of the term), but our word “music” descends from the Greek work *mousike*, which did *not* mean what we now mean by “music”, though it included in its extension the things we call music. Why would it be more plausible to think that aboriginals playing the *didgeridoo* don’t have music “in our sense”, while the ancient Greeks had music, though they didn’t use *our* concept? They have language and depiction *in our sense* (though different from *our* language and *our* depictions). Why should music and art be any different?

I’d like to propose an argument analogous to the “indiscernible duplicates” argument the proceduralist uses to make the point that the definitional properties of art are relational. By analogy, we can hypothesize that even though “western” instrumental music without title or program (what is sometimes called “absolute music”) is a particular cultural phenomenon, it “reveals” the nature of music by setting aside everything which is non-essential to the core of music: e.g., the representational content of sung texts, the ritual functions of music in specific contexts, etc. Focusing on the nature of purely instrumental music (or sung music without words), is more likely to give us a common matrix to all forms of musical expression, both purely instrumental music and “impure” music with sung words. In this way, our “definition [of music] should capture the features of sung words that would be absent if those words were merely spoken.”²²²

By putting aside such properties, the features of rhythm, melody, harmony and timbre stand out as the boundaries that define musical experience – the common core that is then inflected or modulated by the addition of representational texts, social functions such as those involved in religious ritual, etc. the same way that depiction is culturally modulated into specific forms, such as *impressionistic depiction*. The variety of culturally inflected depictions constitutes no objection to the idea that there is a common *core* to all depiction, one that “perceptual theories” identify as the experience of *seeing-in* – though different “perception theories” explain this core experience in different terms. Thus, Steven Mithen has observed:

²²² Kania, A. “Definition”, in Kania & Gracyk (eds.) (2011: p. 10).

“Bruno Nettl (...) has written that ‘the only even remotely reliable guide to the origins of music is the plethora of contemporary and recent musics known to us’. I disagree: while acknowledging the nature and especially the diversity of such musics is important, they are nevertheless the product of more than a hundred thousand years of cultural evolution since the first music arose (...) Indeed, one of my arguments will be that music is relatively unconstrained; once the capacity for making music had evolved, it could be expressed in a great many ways for a great number of different reasons.”²²³

As I see it, an explanation of the nature of music in terms of a core kind of experience, one which is defined along the four basic features of rhythm, melody, harmony and timbre, and which we have described as “the experience of sounds as tones”, is applicable cross-culturally. But suppose that there was some kind of non-western sound-oriented activity to which this theory could not apply... Well, I just wonder why it should matter. Would this mean that music is “a heterogeneous reality”? But we already *know* that – musical traditions differ: some have notational systems, some don’t; some develop one of music’s basic features more than others (e.g. harmony in western music; rhythm in African and Middle Eastern music); some are more centered on improvisation (e.g. jazz); some are more centered on performable *works* (e.g. western classical music). In sum, either music is *one* heterogeneous reality, and in the worst case scenario our definition of music as the art of tonal sound would be adequate to only some of the several world “musics”, leaving us then the burden of explaining why there is a plurality of “musics” that don’t fit the definition – what makes them “musics” instead of something even more “heterogeneous”; or the “heterogeneity” is not at all within one reality (the traditions that don’t fit the definition are not “musics” at all), in which case we need not bother about “universal validity”.

Does this mean that *nothing* could count as a counterexample to our definition? No. There are two major classes of purported counterexamples to the definition of music as the art of tonal sound. Namely, examples of things that are music but not art, and examples of things that are widely considered music but exhibit no “basic musical features”. We will approach these examples, and how our definition may deal with them, in the next section, where I survey some extant definitions of music in the recent literature.

²²³ Mithen (2004: p. 282).

5 Extant definitions of music: a survey

5.1 Jerrold Levinson: music as the sonic intensification of experience

“Music =_{df} sounds temporally organized by a person for the purpose of enriching or intensifying experience through active engagement (e.g., listening, dancing, performing) with the sounds regarded primarily, or in significant measure, as sounds.”²²⁴

In “The Concept of Music”, Chapter 11 of his book, *Music Art and Metaphysics*, Jerrold Levinson makes one of the rare attempts in recent literature to present a definition of music, and the result of his efforts is the above quoted definition. Unlike the theory he has developed about art, it is a kind of functionalist definition. He thinks of music in terms of the function of “enriching or intensifying experience through active engagement” with sounds *as sounds* (as opposed to, say, sounds as speech), where there are several modes of “engaging actively” with sounds as sounds (as in *performing* the sounds; *listening* to the sounds; *dancing* to the sounds, etc.).

In a footnote to the text (p. 274), Levinson discusses the possibility of defining music historically, replicating the structure of his definition of art: “X is music *iff* it is humanly and temporally organized sounds intended for overall taking as any prior music is or was properly taken.” He argues that while certain transformations in the history of music might make that move necessary, we have not yet reached a stage where “music’s essence, like that of art generally, had become almost purely historical, all traces of traditionally musical intent having been finally extruded.” This observation is important for more than one reason, for this amounts to a recognition that in the absence of those developments in the artworld that motivate the historical definition of art, art is to be accounted for in terms of functional properties and the intentions of imposing the relevant functions. If this is a recognition that art must *begin* by being functional, Levinson is basically agreeing with Stephen Davies concerning the first stages of art making: these objects are not yet *conceptualized* as art, neither are they seen as belonging to an institutionalized division of art behaviors. Art must emerge as the purposeful exploration of aesthetic qualities in objects, independently of what social structures are in place. It is important to note that Levinson’s historical

²²⁴ Levinson (1990: p. 273).

definition of art is phrased in terms of *intentions* and *historical connections* between intentions. The crucial concept is that of making an object with the intention that it be regarded as previous artworks are or were correctly regarded. This intention can be either *transparent* or *opaque*, that is, it can be an explicit, direct intention that the object be regarded as certain *particular* previous artworks are or were regarded; or it can be an intention that the object be regarded in a way that *happens to be* a regard that previous artworks were correctly accorded. One crucial difference between this approach and aesthetic functionalism is that the *content* of the relevant intentions has been historically relativized. There are no a priori constraints on what can count as a regard-as-a-work-of-art. Historical development of the arts leads to an increasing diversification of the ways in which artworks are correctly regarded, and thus to a corresponding diversification in the content of the intentions relevant for art-status. This is what Levinson points out that *can* happen in the case of music, though “we are not *quite* at that stage yet.”²²⁵

This means that, at least for the moment, there is one recognizable “musical intent”, whose content is not unconstrained or historically relativized, and it is in terms of that intention in the production of a sound event that music is to be defined. This is the intention of “enriching or intensifying experience” through active engagement with the sounds “regarded primarily, or in significant measure *as sounds*”. The added clause on temporal organization specifies that besides being made with the relevant intention, the sound events produced must exhibit a minimum of “temporal structure”, and in engaging with sounds *as sounds*, their temporal structure must be one of the properties with which we engage.

One noticeable aspect of this approach is its closeness to an aesthetic definition, since “enriching and intensifying experience” seems to be the functional point of intentionally endowing something with certain aesthetic qualities. Arguably, Levinson eschews this option because the concept of aesthetic experience has fallen into discredit after Dickie’s criticisms of that notion, in the course of his well-known debates with Monroe Beardsley in a series of articles, particularly the notorious “The Myth of the Aesthetic Attitude” (1964). This called into doubt the idea that there is a phenomenally or epistemically distinct kind of experience which we call “aesthetic experience”, so that a definition of art in terms of this kind of experience would inherit the vacuousness of the notion, failing to specify anything at all. Levinson avoids such issues by talking about “intensification of experience”, without reference to the aesthetic. However, this ignores the possibility that the exploration of aesthetic qualities doesn’t require any distinct kind of *experience*.

²²⁵ *Ibid.*, p. 274.

After all, we experience a diverse range of perceptual properties of objects without there being the need to specify different kinds of experience corresponding to each kind of perceptual properties. Analogously, our experience of artworks may involve an aesthetic component without it being the case that this component individuates a distinctive kind of experience. We can acknowledge that any experiences of reasonable complexity have an aesthetic component (in the sense that they involve awareness of aesthetic qualities) without grouping the experience of aesthetic qualities into a phenomenally distinct kind of experience. For instance, there is no need to postulate a phenomenally distinct *aesthetic pleasure* – a pleasure that “feels” aesthetic, by contrast with other pleasures. Pleasure is pleasure, and there is no relevant difference between the pleasures of reading, of observing a painting, of listening to music, of playing chess or smelling coffee; except, of course the difference in the *objects* of such pleasures. This allows for a significant distinction between making something with the intention that it has certain aesthetic qualities and making it with the intention of generating aesthetic experiences. Experiences of artworks and modes of attending to them are quite diverse, but they all will involve awareness of aesthetic qualities. Maybe this confusion is due to the fact that aesthetic experience was traditionally conceived as a kind of *peak* experience, whereas all everyday experiences may involve an aesthetic component. Our awareness of aesthetic qualities is not limited to when we appreciate things “for their own sake”, when we are in a special frame of mind, adopt a special kind of attention, or obtain a special kind of pleasure from the experience. The awareness of qualities such as elegance, graciousness, delicacy, boldness, beauty, etc. introduces no difference of kind in our experiences, the same way that awareness of non-aesthetic perceptual qualities (redness, blueness, squareness, etc.) requires no difference in kind in our experiences (e.g. color experiences are just garden variety experiences with a color component).

Levinson’s definition has received some criticisms. Namely, both Stephen Davies and Andrew Kania have pointed out that Levinson’s definition excludes *too much* from the extension of music, on account that music is not always intended to “call attention to itself” or to “enrich or intensify” the experience of those who perform and listen to it. Davies gives the example of *Tafelmusic* (table music) – a kind of baroque composition whose purpose was to serve as accompaniment to aristocratic meals. Davies suggests that “it is possible that this was sometimes not even registered as music by those engaged in eating and conversation. A composer who wrote table music to draw attention to itself simply did not understand such music’s function, which is to be self-effacing.”²²⁶

²²⁶ Davies (2012a: pp. 536-537).

Other examples given are those of film music, which is mainly intended “only for liminal, perhaps unconscious, awareness on the part of the engrossed viewer”;²²⁷ and songs such as “Happy Birthday”.

I believe these objections can be met if we bear in mind that the presence of aesthetic qualities in sounds is compatible with the experience of those sounds not corresponding to a kind of “peak” experience, in which the listener engages the music in a contemplative way, or an exclusive aesthetic focus on the appearance of sounds. Nonetheless, the experience of listeners is “enriched” or “intensified”: try watching a movie with and without its soundtrack and compare the “intensity” of the experience; or try and commemorate someone’s birthday by merely speaking the words of “Happy Birthday”. The music in such examples is there serving the function specified by Levinson, though not calling attention to itself. Furthermore, these examples introduce complications in that the music is conscripted to perform its function as part of an experience which is not an exclusively musical experience. Film music is to be listened as part of a more complex art object, as is the case with opera and theatre music. Of course, in some cases the music will be more “self-effacing” than in others, but this doesn’t mean that the character of the experience is not (positively) inflected by the music. Arguably, in all the cases mentioned there is some level of engagement with music. Levinson seems aware of this when he says that

“there are musics in the world which do not seem aimed at what we can comfortably call ‘aesthetic appreciation’. Music for the accompaniment of ritual, music for the intensification of warlike spirit, and music for dancing are all examples of musics whose proper appreciation does not involve contemplative and distanced apprehension of pure patterns of sound, or put otherwise, does not call for specific attention to its beauty or other aesthetic qualities.”²²⁸

It seems clear that Levinson thinks these musics are designed for some kind of engagement with the sounds, though they fall short of requiring an “aesthetic attitude” as this concept has been worked out by modern aestheticians. Ritual music, music for the “intensification of the warlike spirit”, dancing music, etc., all specify modes of engagement with the sounds *as sounds*: the stirring of emotions through the music (music whose primary purpose is not that of “enriching” or “intensifying” the experience of those who engage with it *in the same way* that concert music “enriches” or “intensifies” the experience of those who attend concerts) is not the same as stirring

²²⁷ *Ibid.* p. 537.

²²⁸ Levinson (1990: p. 272).

the emotions by way of incendiary *speeches*. People who engage martial music in a military parade, reacting to it the way it was designed to; and people who engage with religious music in the way it was primarily designed to, are not engaging with the sounds as vehicles of linguistic meaning or any other extra-musical content, but neither are they engaging it the way a concert attendant would engage it (e.g. a non-religious person listening to a performance of Bach's *Mass in B minor*.)

To replace the "aesthetic idiom" with the idiom of "intensification of experience" may reflect a pervasive (and misleading) association between the aesthetic component of an experience and a particular *mode of attention* or *mode of attending* to the objects that exhibit a given aesthetic profile. This identification of the aesthetic with a mode of attending is perhaps rooted in the way that Kant's ideas about the *disinterested* nature of aesthetic *pleasure* have been later associated with a disinterested "frame of mind" as a condition to the experience of beauty (aestheticians in Kant's time were not concerned with other aesthetic qualities beyond the beautiful and the sublime). This was an important factor in generating the idea that "aesthetic experience" is a special kind of experience, in contrast with the more plausible idea that all kinds of different experiences have an aesthetic component, just as they have other components, without thereby forming a special kind of experience. I can be aware of both positive and negative aesthetic features of objects without thereby having a kind of "peak" experience, characterized by a special "mode of attention". I see that the tea cup has an elegant design, just as I see that the tea cup is white and has a round shape. Awareness of these qualities may involve some sort of hedonic response (of pleasure or displeasure), but there is no threshold in the intensity of hedonic response, below which I can no longer experience the aesthetic quality of the object. Furthermore, Kant's position can be construed in a way that does not confound the nature of the *pleasure* we typically have in (positive) aesthetic qualities²²⁹ with the nature of the experiences of artworks. For one, Kant's theory of the disinterestedness of aesthetic pleasure is focused on experiences of the *natural world*, which afford what he termed "free beauty" (that which "pleases universally without concept": in other words, you don't have to think, say, of a flower as being the plant's genitalia to experience it as beautiful). When it comes to artworks, Kant deploys a concept that seems in tension with his former concept of "free beauty", which is that of "dependent beauty", the beauty that something has *as a certain kind of thing* (e.g. a building may be beautiful as a library but not as a church.) Artworks are human artifacts and as such they differ relevantly from natural objects. To the extent that appreciation of

²²⁹ It is important to remember that for Kant there is no such thing as aesthetic "qualities" in the sense of real properties of objects. To experience an object as beautiful, in Kantian terms, is to have a certain kind of *imaginative* experience, one that is inflected by an "interplay" of the "cognitive faculties".

artworks involves concepts of the object's function, experiencing them as beautiful may involve entertaining thoughts about the purpose with which they were made. Now, I'm sympathetic to an interpretation of disinterestedness according to which the disinterested nature of the *pleasure* taken in something's aesthetic profile does not impose a specific mode of attention to the object. It merely specifies that to experience something as beautiful is not to experience it as satisfying some occurrent desire we have (as when one says that such and such is "music to one's ears", while not meaning it as a literal statement). In this way, the experience that the religious worshiper has of religious music has an element of "disinterested pleasure" because the beauty he hears in the music is independent of his occurrent desires, though his *experience*, on the whole, is "interested" in the sense that his interest in the music (the value that the music has for him) is primarily in how the music will aid the establishment of some kind of devotional state. But this extra-musical function of the music is compatible with the aesthetic profile of the music being essential to its attainment, the same way that a clock performs its function of generating beliefs about the time by way of its other properties. The idea is that for religious, martial or dancing music to perform its function, it must perform it *aesthetically*, if at all. Delicate, gracious and mellow music will not work as martial music; and powerful marches will not be suited to aid in the establishment of a devotional state. But experiencing the music as having these qualities is not dependent on the previous desires of those who experience the music with an extra-musical purpose in mind. Otherwise, how could one tell when some piece of music is *not* adequate to its intended function? Furthermore, this is what makes it possible for a non-religious person to appreciate religious music. I don't think that her appreciation of the music must involve abstracting from the music's devotional function. To an extent, appreciating such music "for its own sake" may involve capturing (making one's own) something of the phenomenology of the religious listener, even if this engagement falls short of a commitment to religious beliefs.

There have been some objections to the idea that the "aesthetic experience" (conceived narrowly as a kind of "detached", contemplative experience) is the only legitimate mode of attending to artworks. Namely, Nicholas Wolterstorff has attacked what he describes as the idea of "art coming into its own" in the culture of the art museum and the concert hall.²³⁰ He defends, precisely, that other modes of engagement with artworks, such as the religious or devotional one, are equally legitimate. One of the examples he provides is Maya Lin's well-known memorial "to those who fell in the Vietnam War". This artwork has a primarily *memorial* function and to attend it or engage it in

²³⁰ See, in particular, his 2003 article "Why Philosophy of Art Can't Handle Kissing, Touching and Crying", *The Journal of Aesthetics and Art Criticism*, Vol. 61, n° 1, pp. 18-27.

a “detached” manner is to miss its point. I partly agree and disagree with Wolterstorff here. I agree that not all art requires (and it is doubtful that any art does) the kind of ceremonial attitude that the culture of the art museum and the concert hall supposedly demands of its attendants (an attitude which is no less “ceremonial” than religious attendance to art). At least some art requires a different approach, but that doesn’t mean that aesthetic profile is not a crucial ingredient of the required mode of engagement. A pink replica of Lin’s Memorial would not be adequate to its function and it would fail to elicit the kind of response required of its viewers; and this is because a pink replica would have a wholly different aesthetic profile. So, one must distinguish between the aesthetic profile and the mode of attending that a particular artwork with that profile requires. The aesthetic profile is of instrumental value in eliciting the latter, even if this isn’t achieved by making the viewer or listener focus (or focus exclusively) on the object’s aesthetic profile; this may only be an essential background to the required response or mode of engagement, yet no less essential because it is *background*. But I can see that Lin’s Memorial has the aesthetic profile it has, independently of any interest I may have in the monument’s representational properties and its memorial function. Otherwise, nothing could be wrong with the pink replica.

Such functional multivalence of art seems to block the idea of one single appropriate way of approaching artworks *qua* artworks; and music exemplifies such functional multivalence to an extreme degree, which suggests that attending to sounds *as sounds* in a way that serves the purpose of intensifying human experience actually fits a wide diversity of “modes of attention” or appropriate ways of approaching music, which determines that different surface properties of music will be more or less emphasized in different traditions where a different “mode of attention” prevails. As Levinson notes:

“Different musics are important to us for different reasons. Some music enriches primarily through conveying emotional content, some more through stimulating the ordering faculties of the mind. Some music exercises the imagination in virtue of its representational aspects, some music satisfies particularly by giving a sense of tradition and shared community. Listening to some music seems to provide insights into the mysteries of the psyche; listening to other music seems to illuminate the underlying pattern of the universe at large. Some music makes one pleasurable aware of one’s body; other music makes one feel bodiless. And some music seems capable of effecting real and relatively permanent alterations in the outlooks and attitudes of listeners.”²³¹

²³¹ Levinson (1990: p. 278).

Experiences of music are variegated (as the above quote makes explicit), exhibiting a wide variety of “cultural inflection”, and this feature tracks the relative value that some surface properties of the music have to this or that non-musical function. Levinson’s specification of the “musical intent” in terms of the “intensification of experience” rather than in terms of generating aesthetic properties or aesthetic experiences is meant to make the definition as inclusive as possible (and also amenable to cross-cultural application) without making it redundant (or unilluminating, as a purely extensional definition would be). But an account of music and musical experience in terms of the aesthetic will only be restrictive if we associate the aesthetic with a specific mode of attending to artworks in general and music in particular. We can think of the “intensification of experience” in terms of the *purely aural* experience, as the “concert hall culture” seems to foreground, but musical experience will typically be “richer”, given music’s multivalence, its ability to serve a variety of purposes. But we can identify a “core” in all this variety of experiences, one in virtue of which the experiences are all *musical*, as well as being culturally inflected in an indefinite number of ways, depending on *other* values that the music is meant to subserve and which of its surface properties are more suited to each purpose. This is true of Levinson’s several examples of “functional” music. Different aesthetic profiles are more or less adequate to one or other function that goes beyond the core of musical experience.

We can now articulate the following question on Levinson’s definition: in virtue of which does music “enrich” or “intensify” the listener’s experience? This will be either in virtue of the music’s aesthetic profile or in virtue of the music’s other properties. If it is in virtue of other properties, one may then ask whether these properties are relevant for the sound event’s status as *musical*. If only aesthetic qualities are relevant for that status, then Levinson’s definition will collapse into an aesthetic definition of music, but many of the criticisms raised against such reading can, I think, be met in the way I’ve tried to meet them. I’m not sure whether Levinson’s qualification that the engagement with the sounds as sounds should be an “active” engagement will block the kind of reply I’ve sketched to the critics of his definition. Are people who watch movies engaging “actively” with the soundtrack? Intuitions here may be divergent.

Other noticeable aspect of Levinson’s account is that he explicitly excludes Muzak as a case of music. Here his argument that the sounds are not designed to be attended as “experience-enhancers” is somewhat more persuasive, but the result is no less counterintuitive: we are inclined to classify Muzak as music, even though it is *bad* music, “put toward a mercenary purpose”, in Kania’s words (to increase people’s propensity to buy things in a certain environment). This casts some doubt on whether we should endorse the condition that in order to be music sounds must be designed to be attended to. Levinson places Muzak on a par with a thought-experiment case where sound events

are devised by psychological researchers in order to make subjects previously put in a “semiconscious condition” enter “psychedelic states of marked pleasurability.”²³² This would not count as music because it fails to satisfy the condition of being meant for active engagement. These cases do not seem amenable to a response in the lines that I’ve suggested – that “background music” is still the object of *some* degree of attention or engagement. Kania will argue that there are many situations where one can use music in a way that requires no engagement whatsoever with the particular qualities of the music. One example he provides is that of his playing the opening of the first-violin part of Strauss’s *Don Juan* to startle a friend awake. This objection might be defused by arguing that the sounds *were* in fact designed to engage attention, but are here conscripted to perform a different function. Analogously, that a baseball bat can be used as an assault weapon does not show that it isn’t a baseball bat. It *is* a baseball bat because it has a certain physical structure, and its history of production traces back to certain intentions and insights, regarding the physical structure and the performance of the intended function. As it happens, its physical structure makes it adequate to be used as an assault weapon. Likewise, any sound event with certain acoustic properties is adequate to startle someone awake. What makes those sounds *music* is a different matter.

A more troubling counterexample Kania gives is that of someone practicing scales and arpeggios on a musical instrument. These seem to constitute examples of *music* which neither call for attention or engagement from listeners, nor are they cases of *art*. I’ll approach these examples while discussing Kania’s proposed definition of music.

Yet another cause of concern is the apparent commitment that Levinson’s definition puts on him regarding a potential distinction between a musical and a non-musical “art of sound”. If there is such a thing as a non-musical art of sound (apart from spoken poetry, prose and drama), it is arguably made with the purpose that people will “engage actively” with it. But Levinson’s approach does not provide us with the tools to make this distinction. According to his definition, any art form that takes sound as its medium and aims at “enriching” or “intensifying” the experience of the listener is music. Another potential difficulty lies in the clause that musical sounds are to be heard *as sounds*. Now, this specifies the mere acousmatic stance, in which sounds are heard in abstraction of their causes, that is, as sound events “detached” from the other physical events that produced them, appreciated exclusively for their surface properties. This implies that for Levinson, for

²³² Levinson (1990: p. 273).

instance, compositions of *musique concrète* are music *in the same sense* that a Haydn string quartet is music, that is, both examples are music *in virtue of the same properties* (being made with the relevant intentions and exhibiting audibility and temporal structure). Levinson's is thus what Andy Hamilton calls a *universalist* position: music is the only art of sound. I discuss this topic in the next section, where I survey Andy Hamilton's view.

5.2 Andy Hamilton: music and (non-musical) sound art

“I defend the claim (...) that music is an art at least with a lower-case ‘a’ – a practice involving skill or craft whose ends are essentially aesthetic, that especially rewards aesthetic attention – whose material is sounds exhibiting tonal organization.”²³³

The first qualification to make while approaching Andy Hamilton’s views on this subject is that he does not propose the above as a *definition* of music but as a set of “salient features” by which we can identify central cases of the music-making activity. In fact, it would not be difficult to come up with some counterexamples if it was being proposed as a definition, since one can imagine cases where music is not conceived to “especially reward aesthetic attention”, and so forth. Nevertheless, there are some points of affinity with the definition I have proposed from the beginning of this work, since Hamilton focuses on features such as the preponderance of tonal material as characteristic of music, though a musical piece can also incorporate all kinds of sounds into its structure and may exhibit a significant *lack* of tonal organization. The definition I have proposed relies heavily on the concept of the aesthetic, though such definition is compatible with music not always being used with the purpose of rewarding aesthetic appreciation. This is because there is a difference between an object’s having an aesthetic profile and its *appreciation* focusing especially on the qualities that constitute that profile. It may be the case that for something to be music it must have *some* aesthetic profile, without it being the case that it was designed with the primary aim (if at all) of rewarding aesthetic appreciation. Thus, for instance, one counterexample that was given against Levinson’s definition of music – “Happy Birthday” – can be said to have an aesthetic profile, yet it wasn’t at all designed for “appreciation”, but to serve a rather different purpose. Of course, given the way I treat the concept of an aesthetic quality it is hard to think of a genuine example of something that *doesn’t* have an aesthetic profile (aesthetic qualities are as ubiquitous as color properties in objects of visual experience). Practically all of our conscious states will have an aesthetic component – the aesthetic components of experience are but a consequence of our “contextual embeddedness” as cognitive agents, our capacity to form desires as well as beliefs, besides having a perceptual apparatus. We see things as endowed with qualities of attractiveness, repulsiveness or blandness, in virtue of how they stand to our cognitive abilities, imagination and

²³³ Hamilton (2007: p. 40).

desires. What distinguishes most art objects is not that they, as opposed to “mere real things”, afford us “aesthetic experience” (this would be like the nonsensical claim that paintings afford us “visual experience” to a higher degree than “mere real things”), but that they are *made* in a conscious, purposeful effort to *explore* aesthetic qualities and our ability to experience things aesthetically, and this can be made to serve the aims of “aesthetic appreciation” as well as some further aim, like that of religious ritual. We can listen to religious music with no other aim than appreciating its aesthetic qualities (which may involve a “sympathetic stance” with the phenomenology of the religious worshiper, not unlike what is required from us in order to appreciate fictional stories), or we can engage with it also for the purpose it was designed, that is, as an aid in bringing about devotional states. In *both* cases, I claim, the aesthetic profile of the music plays an irreplaceable role. If a religious musical piece is serene, restful, delicate, balanced and unified (or powerful, frenzied, and turbulent, such as that of Sufi rituals), those properties are not dependent on the listener’s attitude being one of actual devotion or primarily an interest in those properties themselves; neither are they dependent on the composer having designed the music primarily for one or other mode of engagement. The presence of the qualities is required in both cases.

Now, this might seem to make the idea that music is essentially aesthetic a vacuous one: it would be vacuously true that musical sound events have an aesthetic profile, on account that *everything* (at least that which can be the object of perceptual experiences) has. But the claim is somewhat more demanding than this. All sound events will have some aesthetic profile: they will vary along a continuum from the obnoxious to the enthralling; but not all aesthetic properties that depend on sounds are music-making properties. For instance, a sound may be *delicate* or *eerie* without it being musical. The aesthetic properties that pertain to a sound event as a *musical* sound event are those that depend on those properties of the sounds in virtue of which we hear “movement in an imaginary space”, as specified by Scruton’s theory, regardless of whether or not the *linguistic description* of those properties involves a “necessary metaphor”. These are the properties in virtue of which sounds are heard as *tones*, that is, sounds with “tonal implications” – as exerting attraction and repulse over other sounds, as occupying an imaginary space and travelling through that space. Only the delicateness and eeriness that are instanced by way of such properties (i.e. delicate or eerie *rhythms*, *melodies* or *harmonies*) are cases of *musical* delicacy and eeriness, though non-musical sounds can also be delicate and eerie. The instancing of musical and non-musical “sonic delicacy” is analogous to the way that the very same property can be exemplified in distinct *visual* art forms – e.g. the “base properties” for delicacy will differ in painting and in sculpture.

Hamilton’s most important contribution to the present debate is perhaps that of raising the issue of *non-musical sound art* and its relation to musical sound art. Among the arts whose medium is

sound one can find music, spoken poetry and prose, and drama. But the latter are *impure* in Levinson's sense that the sounds are not to be primarily attended to *as sounds*, but as vehicles of semantic or linguistic meaning. Of course, we also appreciate the aesthetic qualities of words, as instanced in sounds (the elements of speech known as *prosody*); and it is not clear that words *have* any aesthetic qualities independently of their being expressed in sounds. The sadness of a melody is an aesthetic quality of the melody. But the word "sadness" is not sad, and a sequence of words can *sound* sad without sadness being part of the propositional content of what is said. Conversely, a sequence of words can sound joyful even if sadness *is* part of the propositional content. (Word *tokens* seem to have aesthetic qualities because they have *perceptual* qualities; but it is not clear what aesthetic qualities word *types* could have.) When we appreciate such qualities of spoken words we are exploring a music-like property of speech – music also has some language-like properties, like being made of discrete, repeatable elements that perform different functions, depending on the context – this is what motivates the common metaphor of "musical syntax", a compelling though unilluminating metaphor, since music, though specifiable in a notation, is not compositional in the way language is.

Calling attention to the music-like properties of other arts that use sound as their medium is important to Hamilton's view that music stands "on a continuum with non-musical sound-arts, differing from them in the preponderance of tonal material."²³⁴ Any sound can be fused into musical composition, but the predominance of tonal organization is a crucial feature that distinguishes music from non-music, with a range of fuzzy or borderline cases on that continuum. On this view, the concept of music is vague or fuzzy "around the edges", but that doesn't mean it has not a clearly distinguishable core. There are clear cases of music and of non-music, even if there are also some more problematic unclear cases. It is to be expected if music and non-musical sound art stand in a continuum. The situation is not the same as if we had an arbitrary concept. At most, the existence of disputed cases may lean us toward a simple disjunctive definition that specifies both the undisputed center of the concept and a way of dealing with the disputed examples.

Hamilton's characterization of music in terms of "an art at least with a lower-case 'a' – a practice involving skill or craft" is connected with his critique of Roger Scruton's "conservative universalist" position, that is, the idea that music is the only "high" art of sound, or the only art of sound with a potential for "high art" status. Hamilton distinguishes the following positions: a) conservative universalism; b) *avant-garde* universalism; and c) non-universalism. The first two

²³⁴ Hamilton (2007: p. 46).

positions agree that music is the only art of sound (*as sound*, by contrast with sound as a linguistic medium) but disagree on the status of “avant-garde” works that exploit the *acoustic* properties of sound²³⁵ (including what conservatives would call *noise*) in detriment of tonal organization; and the latter position distinguishes, among the arts that use sound (*as sound*) as their medium, both a musical and a non-musical sound-art. Hamilton is sympathetic with (c), arguing that there is perhaps a place for an “art phonography”, paralleling “art photography”²³⁶. It may very well be that we will come to use the expression “aural arts” analogously to the way we use the expression “visual arts”, namely, to discriminate between more than one art form.

Levinson’s definition, as we’ve seen, places him in (b), since he has no conceptual means to distinguish between the “musical intent” of a musician working mostly with tonal material, and the “musical intent” of the avant-garde composer experimenting with “noise”. As long as it has temporal structure and it is intentionally designed to “intensify” experience through the engagement with sounds *as sounds*, it is music. Levinson’s definition can’t distinguish any *non-musical* art of sound *as sound*. If its aims are artistic and if its medium is sound *as sound*, then it is music.

Hamilton’s characterization of music in terms of skilled practice or “craft” is also connected with his rejection of Scruton’s version of the acousmatic thesis, particularly, the idea that only acousmatic components of an aural experience (those that abstract from the sounds’ physical sources) are genuinely *musical*. By contrast, Hamilton presents a “twofold” thesis, according to which musical experience has both acousmatic and non-acousmatic components; that is, in attending to music, we attend not only to an “order of tones” which we experience *in* sounds, but also to the physical properties of sounds as such, as well as the way in which they are produced by musicians. Hamilton is sensitive to the role of *performance* in music, and he thinks that the acousmatic thesis, as proposed by Scruton, unduly excludes properties relative to the performance from the core of musical experience. The idea is that, in some cases, we do not fully appreciate the music if we are not aware of its performative aspects, such as, for instance, the *virtuosity* involved in the performing of certain passages, the strain that goes into the physical action of producing the sounds, which also becomes an audible element of the music itself.

I believe that this and other objections to the acousmatic thesis tend to muddle the waters between epistemological questions about the nature of music appreciation and metaphysical questions about the nature of music and of musical experience. Scruton’s analogy with the case of a figurative painting by Van Gogh proves useful here. Someone who doesn’t see the marks on the

²³⁵ Those properties which are non-acousmatic, in Scruton’s (not Schaeffer’s) sense of “acousmatic”.

²³⁶ *Ibid.* p. 40.

surface of the picture as brushstrokes or that fails to take account of the brushstrokes (maybe because he is observing a poor reproduction instead of the original) will, to that extent, have an impoverished understanding of the painting. (John Dilworth has also argued that a picture can be interpreted in two distinct ways, one correct, the other incorrect, respectively, as a naturalist representation of a hideously disfigured subject, or as an expressive representation of a normal subject; the image itself does not compel one interpretation or other, but we *know* how to interpret it correctly because we can't help taking account of the brushstrokes as a stylistic component.) But impoverished though his experience may be, someone who doesn't take account of the brushstrokes will still have the more basic experience of *seeing* the depicted subject *in* the marked surface. The correct interpretation of the stylistic components is an extra feature, one that further inflects the experience of the marked surface *as* a figurative representation; it is not part of what makes the experience an experience of something *as* a figurative representation. The analogy holds with the musical case: someone who fails to take account of the effort or strain that goes into the performance of a piece of music, or who doesn't know what instruments are being played, may have an impoverished understanding of what he is experiencing, but his experience will not be any less *musical* because of it. Awareness of the performative components is not what makes the experience a *musical* experience. The acousmatic thesis, like Wollheim's account of *seeing-in*, is a metaphysical thesis about the nature of musical experiences; but the objections commonly raised to the acousmatic thesis pertain to the epistemology of musical experience, not the properties in virtue of which the experience is a musical one. If the person observing the Van Gogh painting sees only brushstrokes in the marked surface, while failing to see the depicted subject *in* that set of colored marks, he will not have the relevant experience at all; he will be in a parallel position to that of a listener suffering from amusia, who experiences only noise where a normal listener experiences rhythms and melodic development.

While the visual appreciation of a performance may increase understanding of the experience of the music, such aspects are not essential to the perception of sounds as music. After all, people who suffer from congenital blindness are able to have fully satisfactory musical experiences and even become performers themselves. Or imagine someone who always experienced music through recordings, never having seen a musical instrument or learned about musical instruments. She would still have *musical* experiences, even if she was unable to identify a single instrument or have any idea of how the sounds are produced in those instruments (though she might guess that some sounds are produced by the striking of strings; others by friction, etc.). Her experiences would be musical because her ignorance of sound production means would not prevent her from hearing rhythmic movement, melodic intervals, chords, and the distribution of timbres as organizing the

“virtual” or imaginary space of the music. To claim otherwise would be a mistake analogous to claiming that a normal listener cannot enjoy a purely electronic composition unless he knows exactly how the sounds he hears are produced. In some of these cases, obscuring the source of the sounds is part of the point. Think of those listeners who hearing for the first time Stravinsky’s *Rite of Spring* are unable to grasp that the folk tune with which the work begins is being played on a bassoon in high register. No one is prevented from hearing the melody (the musical content of the sound sequence) on account of that ignorance, and there is room to doubt whether upon learning how the sound is produced it was one’s *musical* understanding, as such, that was enhanced. Knowing how to tell the time by looking at a clock is one kind of understanding; knowing how a clock works is a distinct understanding. Knowing how a clock works enhances one’s understanding of clocks, but it introduces no change in one’s ability to tell the time. The question is whether musical understanding is relevantly similar to the latter. One can experience the musical properties of a sound event while remaining ignorant of properties pertaining to the performative means by which the sound event is generated: one can form appropriate auditory expectations, to feel the “gravitational pull” exerted by some notes and chords over others, to respond emotionally to music in a way appropriate to the character of the music, etc., while ignoring substantive aspects of the performance. One’s feeling of rhythm is not impaired by ignorance of what percussion instrument is being used to generate that rhythm, and so on. And why stop at information concerning sound sources and performance? Once one has crossed the acousmatic line, why not include details of instrument building and acoustics as essential components of *musical* understanding? Is one’s experience of rhythm, melody and harmony aesthetically inflected by knowledge that certain sounds are being produced by frictionizing synthetic or animal gut strings with a horsehair bow? Is it *musically* relevant whether a certain “woody” sound is actually not produced with a wood instrument? Or is the precise nature of the wood musically relevant?

Of course, the acousmatic thesis could be interpreted as making an epistemological point about the appreciation of music. It could be construed as the implausible claim that musical experience is not at all affected by knowledge of performative means and sound sources. But this would be a misunderstanding. It would be like claiming that knowledge of watchmaking is irrelevant to one’s understanding of clocks. But the issues are confounded: to experience sounds as music is like being able to tell the time by looking at a clock; to understand how clocks work is a different matter. And is one’s aesthetic appreciation of clocks inflected by knowledge of watchmaking? If it is, where does one draw the line? Is knowledge of the metal alloys that go into the constitution of the clock, or detailed chemical knowledge of the enamels used to decorate a clock’s dial plate, aesthetically relevant? Is one’s aesthetic appreciation of a dance affected by detailed knowledge of the human

motor system? Does it affect in any way our ability to perceive a movement as graceful, delicate, elegant or awkward? This is perhaps a better example than the clock, since we have a clear distinction between 1) what enables human beings to move in certain ways, 2) the aesthetic qualities of the movement itself; and it provides a useful parallel with a) musical movement and b) properties of the sounds on which the musical movement depends. The acousmatic thesis asserts that we can experience (a) while ignoring (b).

There is perhaps an illuminating analogy to be made with the experience of aesthetic qualities in general: my experience of something's having a determinate aesthetic quality, such as elegance, grace or garishness, is compatible with my inability to spell out precisely on what non-aesthetic properties of the object those aesthetic properties depend. The ability to give specific and illuminating descriptions of how a particular aesthetic property of an object emerges from its other qualities may enrich one's experience, but it is not what makes one able to experience the aesthetic quality. To claim otherwise would be to reverse the order of explanation. Upon having an aesthetically intense experience of a painting, a musical work, or whatever, one tries to understand what aspects of the experienced object account for the character of the experience. We do not engage in musical analysis in order to acquire the ability to experience the music aesthetically. Rather, the reverse is true: we are motivated to learn more about music because we were previously aesthetically moved by our encounters with it. We want to know how it works because it already holds interest for us. Likewise, one is drawn into aesthetics because one is curious about our experiences of beauty and other aesthetic qualities. One does not delve into aesthetics to *acquire* the ability to experience things as beautiful or ugly, etc. (In fact, if aesthetic predicates were utterly meaningless to us, the wish to acquire the ability to experience things aesthetically would be unintelligible. One *can* wish to be more aesthetically discriminating than he presently is because aesthetic predicates are *not* utterly meaningless to him.)

We have seen that Levinson's definition of music allows for no distinction between musical and non-musical art of sound (even adding the qualifier "as sound", to contrast with the use of sounds as vehicles of semantic meaning). The purpose of enhancing or intensifying experience is wide enough to cover all artistic intents: if an artwork does *not* enhance any aspect of our experience (hedonic, aesthetic, cognitive, or whatever) then what was the point in making it? What reasons would we still have to discriminate a category, no matter how fuzzy around the edges, of *art objects*? Philosophers may disagree about the properties that are relevant for art status, but *some* properties must be relevant, even in an anti-definitional approach. Even the enemy of definitions that takes time to experience artworks does it for a *reason*, because *some* properties of the objects in question are *worthwhile*, because they positively impact one's experience in some way or other. Whether to gain

knowledge of the world, to have “aesthetic experiences” or just for a “good tingle”, there is some aspect of our experience that is enhanced by art, and even when art does not achieve this purpose, it seems that what makes it intelligible as a human practice is the *attempt* to achieve some such purpose. (This is not to be confused with the idea that artists are mainly or exclusively moved by altruistic impulses). Making *bad* art is intelligible as a failed attempt at achieving some valuable artistic purpose; and objects that seem to defeat the purpose of art often draw their force, precisely, from expressing thought-provoking ideas about the purpose of art; which has the effect of making them serve *some* artistic purpose after all. My experience of the least original attempt at conceptual art is still distinguished from my experience of natural objects in virtue of this aspect; which doesn’t imply that the latter experience can’t be *more* rewarding.

But now one may ask: does it *matter* whether we distinguish between music and non-musical sound art? Does it matter whether or not we use the word “music” to refer to all sound events designed to “intensify” our experience? Is there something to be gained by distinguishing further, between musical and non-musical arts of sound? Isn’t this distinction *ad hoc* or arbitrary?

Actually, it isn’t. To claim otherwise would be like questioning the relevance of distinguishing photography from painting or other visual arts. After all, painting, photography, etchings, woodblock prints... They’re all aimed at enhancing our *visual* experience. Why further discriminate between them? The answer is that discriminating is useful *if* there are real differences to account for; if *not* discriminating would make us miss some relevant aspect. And the relevant aspect is not whether we apply this noise – “music” – rather than some other noise, to this instead of that. If the only interesting unifying property of the things we call music were the fact that we apply that noise to them, it would not matter, but neither would it be intelligible for us to have such a category. We feel the need to categorize because we intuitively understand that there *are* differences. And we need some device for marking the difference. Having discerned the difference, it doesn’t ultimately matter whether we use the same noise to encompass the different things in question, as long as we understand each other.

But what *is* then the relevant difference between musical and non-musical arts of sound? As it happens, the difference hinges on a technological development, in the absence of which the existence of more than one art of sound is much more difficult to conceive. According to Hamilton, “recording was the crucial technological advance which liberated a concern with sound as sound, in the sound art sense – never before could one analyze sounds and alter their envelope.”²³⁷ Recording

²³⁷ Hamilton (2007b: p. 47).

made it possible for all sorts of sounds to be abstracted from their circumstances of production and exploited for artistic purposes; it also allowed composers to manipulate sound directly, instead of producing abstract instructions or recipes for the *performance* of sound events. In fact, recording allowed for a kind of composition that eliminates performance altogether – many electro-acoustic or electronic compositions are not even *performable*, they are for *playback* on appropriate decoding devices, not for performance. Unless one listens to a playback from the original or a suitable copy of it, it won't be an experience of *the work*. This opens up the possibility that such pieces, unlike traditional music for performance, are susceptible (in principle at least) of reproductive forgery (e.g. in the form of perceptually indiscernible replicas). In this sense, to the extent that some electro-acoustic compositions count as music, recording seems to challenge Goodman's view that music is an allographic art. Since the advent of recording technology, and to the extent that some sound compositions for playback are music, at least some music is *autographic*. So, whether or not this technology has prompted the emergence of a non-musical sound art (an "art phonography"), it *has* managed to introduce ontological complications in the category of music.

This was the idea behind *musique concrète*: what made this music "concrete", by contrast with the "abstract" music for performance, based on tonal organization, was not merely the use of everyday artificial and natural sounds, but the *direct* manipulation of sounds. Where previously the composer issued instructions for performance, or *blueprints* for the multiple realization of rhythmic and harmonic relations, the "sound researcher" in the tradition of Pierre Schaeffer and Pierre Henry liberates composition from the abstract blueprint and works on concrete, individual sounds, just as the sculptor works on concrete, individual blocks of stone. Recording allows the appreciation of these everyday sounds in abstraction of their physical sources; it allows them to be appreciated for their purely sonic qualities, and not as "acoustic cues" to navigate the environment. The manipulation of recorded sound adds another layer to this expansion of the acousmatic experience, by obscuring the causal origins of sounds and allowing the purposeful exploration of noise elements for their aesthetic qualities.

Hamilton makes some more fine-grained distinctions that might be useful in organizing our "conceptual map" to sound-oriented activities. Non-musical sound art is to be distinguished from the category of *sound design* (with the implication that things in this category are not aimed at artistic purposes). Sound design branches in two main subsections: a) *significant sound-design* (this includes: "mobile telephone ringtones, alarm clock tones, car horns, door chimes, computer 'ear

cons' and airport announcement chimes").²³⁸ These are functional sounds, designed to convey "non-natural meanings";²³⁹ b) *non-significant sound-design* (this includes: "fountains, car or aeroplane engines, and tones for synthesizers and other electronic instruments").²⁴⁰ The category of *sound art* also branches in two main subsections: c) *documentary sound art* (which, as the name suggests, aims at documenting sonic environments, though, as Hamilton notes, "any artistic creation, even if it seems to involve only selection of material, has a form . . . basic if mundane compositional decisions are still required – when to start recording, the location and direction of the microphone, and so on."²⁴¹); d) *non-documentary sonic composition* (which consists in creating and not merely documenting a sonic environment). *Music* is probably best thought of as standing in a continuum with non-documentary sonic composition, with clear cases of both categories on each side, a range of intermediate or hybrid cases, and a much thinner section of borderline cases – things that can fall under one or other category, depending on contextual factors, for instance, that the composer intended it to be listened to as a piece of music, and so, against the background of a musical tradition.

²³⁸ *Ibid.* p. 62.

²³⁹ "Non-natural meaning" (i.e. conventional meaning) is to be contrasted with cases of "natural meaning", e.g., where the presence of smoke indicates the presence of fire, or a tree's rings indicate its age.

²⁴⁰ Hamilton, *ibid.*

²⁴¹ *Ibid.*

5.3 Andrew Kania: artless music and listening for basic musical features

“Music is (1) sounds (2) intentionally produced or organized (3) *either* (a) to have some basic musical feature, such as pitch or rhythm, *or* (b) to be listened to for such features.”²⁴²

Kania’s definition can be viewed as an extension of the definition I have proposed from the outset, namely, that music is the *art of tonal sound*. It can be seen as an extension of it in the sense that while it draws essentially on “basic musical features” (the properties in virtue of which we are able to “hear sounds as tones” – the NKA in terms of which the practice of music making can in turn be defined), it makes room for two phenomena that must be considered as distinct possibilities: 1) cases of *music* that are not cases of *art*; and 2) cases of *musical sound art* that are sonically indistinguishable from *non-musical sound art*. An example of the former is that of someone practicing musical scales and arpeggios; an example of the latter is that of an avant-garde piece with no distinguishable “basic musical features” (e.g. Yoko Ono’s *Toilet Piece*, which consists of an unedited recording of a flushing toilet; or John Cage’s *William’s Mix*, which is a tape composition of sounds “painstakingly spliced together without regard to their basic musical features”²⁴³) but made to be *listened for* such features. (I understand the notion of “listen for” basic musical features as a kind of contrast listening, in which a given sound pattern, either in performance or playback, is to be appreciated against the background of our ability to “hear sounds as tones”.) While I wish to keep an open mind to these possibilities, I don’t want to make any strong commitments to either. Kania’s definition is more inclusive than the definition of music as the art of tonal sound, and this may be desirable to the extent that we must make room for genuine borderline cases, in which the status of something as music may depend not on “natural-kind” features of musical experience but on a culturally arbitrary procedure. Just like a blank canvas may count as a painting within a *tradition* of painting, even if no tradition of painting could be entirely comprised of blank canvases, a musical tradition may include many things that we would not “naturally” classify as music. That is, the procedural aspects that art as a cultural activity (and not merely as a “natural” category, in aesthetic functionalist terms) inevitably integrates in its development may provide a robust sense in which these things are to be thought of as legitimate elements of a musical tradition. But neither is

²⁴² T. Gracyk & A. Kania (eds.) (2011: p. 11).

²⁴³ *Ibid.*, p. 9.

Kania's definition over-inclusive in the sense that it will count as music just anything that happens to be made in the context of a musical tradition or that happens to be made by a musician. For instance, he excludes John Cage's *4'33''*, classifying it as a case of non-musical sound art, since it is not a genuinely "silent piece" and it is not made to be "listened for" basic musical features, but (arguably) to highlight the experience of incidental, ambient noises, which the performance of the piece absorbs in their totality, with the result that no sounds are in fact *ambient* to it: all occurring sounds are to be taken as the piece's "content". But he concedes that a genuinely silent piece, one that is made with the intention that it be listened to for basic musical features, would count as music. He gives as an example his own (thought-experiment) work, *Composition 2009 #3*, consisting of a score that specifies a length of silence framed by the usual gestures that signal the beginning and the end of a performance, adding the clause that the content of the piece is the silence framed, not any of the co-occurring sounds.²⁴⁴

In working out his definition, Kania excludes as implausible some definitional strategies, namely, that of defining musical sounds in terms of purely intrinsic, acoustic properties of sounds (since all manner of sounds can be used to make music, and no range of acoustic properties can be *a priori* excluded from the medium of music), and that of a subjectivist definition, according to which music is whatever is perceived as music by a given listener (or that whatever the listener *decides* to perceive as music, if one can decide over that any more than one can decide to see an object as red or blue). The implausibility of the first strategy was already explained in our discussion of different kinds of functional properties (in section 1.1), as well as in the explanation of what "hearing sounds as tones" consists in (section 4.1). "Basic musical features" are a special kind of relational properties of sounds that depend both on certain intrinsic, acoustic properties of sounds *and* on our natural ability to "hear sounds as tones" (of hearing sounds as abstracted from their causes *and* as exhibiting a special kind of order or organization, one that differs crucially from the grammatical organization of speech). If we lacked this ability, as is plausible that all non-human animals lack, no sounds, however intrinsically constituted, would count as music, not even in the "extended" sense that allows us to classify as music intentionally produced or "organized" sound events that lack "basic musical features". These are only intelligible as music given the existence of a tradition of music-making in the *primitive*, non-derivative sense of producing sound events with the purpose that these exhibit an organization which is not to be confused with the organization of its physical sources nor with the organization of speech. While the "basic musical features" are disjunctively

²⁴⁴ Kania (2010: pp. 350-351).

shared by all central cases of music (some music may be purely rhythmic, or it can be monadic, without a supporting harmony), the wide range of acoustic properties in which these basic musical features are “realized” does not allow a definition in terms of those acoustic properties. Music can be made with out of tune instruments, with instruments of indefinite pitch and with all sorts of devices that were not primarily designed to make music. The crucial features in a musical sound event are connected with the *intentions* of those who produce the sound events that are actually heard (the performers) and those that design the sound events or specify the sound patterns to be performed (the composers). Note that in many circumstances the composer and the performer may be one and the same.

Though it might superficially seem that Kania’s arguments against the subjectivist strategy conflict with his (b) clause – that some sound events count as music if they are to be listened for such properties – this confusion is dispelled by noticing that the definition is producer-centered. In the case of a composition lacking in “basic musical features” that is to be listened for such features, the norms are set by the composer, not the listener. This marks a crucial difference between the case of listening to Yoko Ono’s *Toilet Piece* or to John Cage’s *William’s Mix* (two examples given by Kania of sound pieces without “basic musical features” that are nonetheless to be experienced against a background expectation of hearing tones) and the case where one hears the sound of the clanking wheels of a train as a rhythmic sequence. Many sounds can be heard as *musical* in the sense of being *music-like*, but merely hearing them in such a way doesn’t make them music (this is a distinction also carefully drawn by Levinson). The cases where something counts as music in a derivative sense, by being *about* “basic musical features” while not *exemplifying* them, requires a wide cultural frame, in which the composer is intelligibly recognized as making a contribution to the musical tradition, by producing a piece that “comments” on things that are music in the primitive, non-derivative sense.

Unlike Levinson, Kania considers unnecessary to include a restriction on temporal organization. Levinson excludes from the domain of music the possible case of “colorful instantaneous combinations of sounds” like “chords of vanishingly brief duration” to be “savored independently”. While Kania acknowledges that a tradition of “sonic art” entirely constituted by such cases would (perhaps) not be properly considered a musical tradition, he points out that within a tradition properly called musical there may be individual instances of such pieces. He calls our attention to the fact that we may have an actual counterexample to Levinson’s temporal organization clause in La Monte Young’s *Composition 1960#7*, which consists of a single open fifth with the indication “to be held for a long time”. Though this piece exhibits no obvious temporal structure, we can make it fit the challenge to Levinson by imagining that it is indicated “to be held instantaneously”.

An important aspect of Kania's definition is that it doesn't require that the intentions relevant for the musical status of the sound events produced actually include a concept of *music*. Neither does it require that the producers or designers of the sound events have explicit concepts for basic musical features (such as our concepts for rhythm, melody and harmony), since our ability to hear sounds as tones does not require us to master the use of those concepts (the criteria for the presence of the properties referred to by those concepts is phenomenological). Basic musical features are heard spontaneously and they can also be produced spontaneously. (Musical literacy is not a condition of making music. Analogously, human beings started producing *fictional stories* before they began to grasp grammatical concepts, as well as the concept of *fiction*.) Of course, certain feats of musical intelligence, like counterpoint and intricately complex harmonies, may not be possible without the mastery of the relevant concepts or in the absence of a sophisticated musical notation that allowed the music to become increasingly more complex (both "vertically" and "horizontally", in terms of its texture) for the simple reason that given musical notation, the players don't have to rely on memory for every aspect of the performance, but the point is to emphasize that the definition is not restricted to sophisticated, musically literate traditions. Neither are the sound events to be counted as music only if they are primarily or exclusively conceived to be listened to in terms of rhythm, melody and harmony. Sound events that arise as a consequence of more complex actions, or to serve as accompaniment to dance or ritual (and as such were not conceived to be experienced separately from these actions) still count as music since, to the extent that they are to be aurally experienced (even if not with the supposedly detached attitude of the concert goer), they are to be experienced in terms of basic musical properties (these properties are instrumental to the sounds playing their role in the more complex actions; e.g., dancers must dance *to* a certain rhythm or melodic pattern, however simple or self-effacing).

At this point I'd like to defend Kania's definition from a purported counterexample given by Stephen Davies. According to the latter, the definition is too broad, in that it counts things such as Morse Code as cases of music, since it is to be listened to in terms of rhythm. Another example would be tone-languages and the periodicity of speech, whether or not in languages that use pitch as a significant element or meaning vehicle. Kania dismisses the latter on the grounds that rhythm in music is typically circumscribed to stricter divisions of time.²⁴⁵ This defense is arguably inadequate, as Davies suggests, since not all rhythm in music obeys such strictures – there is the phenomenon of free rhythm and additive rhythm, characteristic of some non-western musics. Davies seems right to

²⁴⁵ *Ibid.*, p. 8.

claim that such a defense would force Kania to count the examples of Morse Code and the periodicity of speech as cases of music. But the objection may be dispelled on other grounds. Scruton's contention that rhythm involves more than periodicity is useful here, but also the point that Morse Code is to be experienced primarily in terms of conventional meaning and the patterns of duration are meaningful only in terms of their contribution to the realization of this conventional meaning. Also, Morse Code is to be experienced in a more "atomistic" fashion: we don't hear a stretch of Morse Code for the qualities of the whole. Rather, we grasp it by the contribution of its individual elements, which are to be taken in themselves, even if we must also take account of the whole pattern to determine the relations between those elements. Altering some bits of that stretch may make it the case that we no longer have a message in Morse Code, but only meaningless noise, while altering notes in a melody or rhythmic sequence does not make it the case that we no longer have a stretch of music. Evidently, the sounds of Morse Code, or something that could be mistaken for it, may be included in a musical piece, or it can be given a *musical* meaning (or role) by being inserted in a musical composition. One case in point is the piece, *Astronomy Domine*, that opens Pink Floyd's first album, of 1967, *The Piper at the Gates of Dawn*, which does not actually contain Morse Code, but something that could be mistaken for it. And the musical role played by those sounds in that context does not require one to try and perceive any "message" or conventional meaning. To hear a sequence as rhythmic is to hear each individual sound as "bringing into existence" the sounds that follow it; it involves detaching the sounds from their sound sources, even if we do this in a single act of awareness, in which we are also aware of the physical production of sounds (the same way we are aware that the picture of the dog is a set of marks on a surface). Just as the clanking of a train can be heard *as* a rhythmic sequence without it actually *being* a rhythmic sequence (by a similar feat to that by which we can hear the ticking of a metronome as grouping of pairs of stressed and unstressed beats, while the sounds themselves do not exemplify this pattern), so we can adopt an attitude of "reduced listening" towards a stretch of Morse Code. But this doesn't make it music. Kania seems to be aware of a reply along these lines when he writes that "While you cannot turn the sounds of a train into music just by hearing them a certain way, you could turn them into music by repeating them with musical intent (as, for example, Arthur Honegger did in *Pacific 231*)."²⁴⁶

²⁴⁶ *Ibid.*, p. 6. Kania's phrasing may be misleading here since it suggests that the *actual sounds of a train* are "repeated with musical intent" in Honegger's work, which only employs orchestral sounds evoking the (more than just aural) experience of a moving train. Obviously, Kania refers here to similarities in the dynamic features of sounds. Thanks to Prof. Vítor Moura for pointing this out.

Likewise, the pitch elements in tone-languages are not to be heard in terms of the contribution they make to a *tonal* organization of the sound events. These are not to be listened to *as sounds* but as “sonic vehicles” for something else – conventional meaning. But in music, what is “conveyed” can only be separated from its “vehicle” conceptually, in the sense that other sound events may exhibit the same organization (the music that is heard in the sound event is *repeatable* or “multiply instantiable”), it cannot be replaced by another medium. What is conveyed by an utterance in a tone-language can be conveyed by a different utterance, but in music the “idea” conveyed is a *sonic* idea or a thought in terms of sound. It cannot be actually detached from the medium of sound, even if it can be abstractly conceived as a *pattern* of sound. Even where a trained musician is able to appreciate a musical piece merely by reading the score, what is actually happening is that the reading of the score elicits an imaginary experience of the sounds, and it is *that* experience which is relevant for the “appreciation”, not the graphic symbols used in musical notation. If no sounds are imagined, no music is heard “in the head” of the musician reading the score.

Another possible counterexample could be that of the role of pitch changes in utterances to mark differences in illocutionary force. For instance, ending a sentence with a rising pitch usually turns the sentence into a question.²⁴⁷ Here too a sound event is produced intentionally to be heard for its pitch qualities, among others, since those qualities are relevant to determine illocutionary force. Once more, what relevantly distinguishes these cases from a genuinely musical case is the kind of functional properties involved in a sound event counting as an assertive, a directive, etc., and the kind of functional properties involved in a sound event counting as a melody. These cases suggest that “being heard for basic music qualities” may be too weak as a condition, to the extent that at least some of these qualities – like pitch and timbre – are shared by music and non-musical entities as structural elements. It is not enough that a sound event should be listened to for its pitch or timbral qualities; we must also know what *role* those qualities are supposed to perform, in order to distinguish cases such as Morse Code and the illocutionary force of sentences from cases where they perform a genuinely musical role.

The purported counterexamples of tone-languages or the periodicity of speech (and the *prosodic* elements of spoken language in general) are useful to further elaborate the distinction between language and music, on the one hand, and, on the other, to see the *continuity* between music and merely music-like phenomena (such as music-like components of speech), without thereby making the distinction between music and non-music empty or inoperative.

²⁴⁷ I took this example from Gracyk (2013).

As pointed out before, a series of utterances may *sound* sad, in a similar way that a melodic line does, without sadness being a part of the utterance's propositional content. This is where language acquires a music-like character that is exploited in spoken poetry and prose, for instance. The relevant difference is that in such cases one cannot perfectly detach the way an utterance exhibits such qualities from the conventional meanings it conveys, even where the propositional content does not include an explicit reference to the properties exemplified by the utterance. In the case of sung words, the propositional content of what is sung can inflect our experience of the melody, but it won't be of any help where the character of the melody conflicts with the character of the words being set to music. A joyful verse set to a sad melody will produce a cognitive dissonance that may have an interesting aesthetic effect (for the good and for the bad), though it won't change the character of the melody. But in a "sad" utterance, the sadness will always be partly a function of what is said, even if the words are not about sadness. On the contrary, the prosodic elements of speech may inflect the meaning of the words, just as a joyful melody inflects the meaning of "cognitively dissonant" words that are set to the music (e.g. making them sound ironic).

Traditionally, in music with sung words or lyrics, speech is made to conform to the musical structures allowed by the instruments. For instance, in opera, the characters "speak" in the tonal intervals allowed by the instrument's tuning system, namely, the twelve semitones of equal temperament. But actual speech, not just in tone-languages but all natural languages, exhibits a much wider range of pitch gradations. In the first half of the twentieth century, the composer Harry Partch had the idea of reversing the model of adapting speech to the musical instruments and devised a wholly new musical system, consisting of 43 (micro)tones per octave (in just intonation), instead of the traditional model of twelve equal semitones per octave. To do this, he had to build several new instruments suited to the task, since music composed in such a system could not be played using the more traditional instruments of western classical music. The works he composed with his new system deploy melodic lines that follow much more closely the patterns of the speaking voice. But the distinction between music and the music-like is still applicable, since the result is not to be listened to in terms of conventional meanings, no more than other, significantly different microtonal experiments. Yet, the fact that spoken language could be used as a model to build a new musical system suggests that the boundary between music and spoken language is vague, in the sense that spoken language, like many other non-musical phenomena, have nonetheless some musical or music-like features that can be foregrounded to produce musical phenomena in the primary, non-derivative sense.

But even radical experiments such as those of Partch and other composers are not only to be experienced in terms of "basic musical features", they patently exhibit or exemplify these features.

So, “radical” as they may be, these are perfectly peaceful examples of music according to the definition of music as the art of tonal sound, and also according to Kania’s definition, as it takes the idea of tonal organization to constitute the undisputed center of what is to count as musical entities. The really “hard” cases are those sonic pieces that completely eschew “basic musical features” but are nonetheless presented as musical compositions. Such pieces, given appropriate contextual changes (such as the composer’s intentions and specifications regarding how the piece is to be correctly experienced), could instead constitute examples of non-musical sound art. These are the ones that truly pose a challenge to the definition of music in terms of tonal organization, i.e., structures of rhythm, melody, harmony and tone color. (Perhaps significantly, Partch was profoundly inimical to the kind of aleatory procedures on which artists like Cage relied to produce their works. He considered that such things could not count as music, since *music* involves the making of compositional *choices* by the composer. As soon as you stop making choices, Partch believed, you stop making music. Whether or not Partch was right (and I incline to believe he was right), this is another example of how a “pragmatic constraint” on the metaphysics of music would not help us to decide the matter, since it is a case where equally “credentialed” musicians disagree in their intuitions over what counts as music.)

Kania’s way of dealing with such borderline cases between non-musical “sound design” and music is to appeal to the relevant contextual factors. For instance, Yoko Ono’s *Toilet Piece* is included in an otherwise uncontentious example of a musical work – her 1971 album *Fly*; and it is to be listened to against a background expectation of experiencing a musical organization in the sounds. By contrast, Cage’s *4’33’’* is not aimed at this kind of “contrast listening”, but more at the complete overturn of our expectations regarding what is to count as “musical”. To the extent that such overturn would make the concept of music empty or arbitrary, the piece fails as a musical piece and it is perhaps more suitably described as non-musical sound art, or even as a kind of dramatic performance piece or a “conceptual work” *about* music. Zappa was wrong: it is not enough that Cage “frames” his gurgling of carrot juice as his composition for this composition to be music; nor is it enough that the audience “decides to perceive it as music”. The status of the composition as music depends on its being meant to be experienced against the relevant background expectation, and *this* expectation entails the traditional distinction between music and non-music in terms of tonal organization (structures of rhythm, melody, harmony and tone color, which are not mere “vehicles” for something that can be conveyed through a non-audible medium).

We can perhaps use a previous notion from our exploration of games and status-functions to shed more light into this phenomenon. Remember those borderline cases of activities that are counted as “games” but are superficially indistinguishable from other activities that would not be

counted as such in “normal” circumstances – such as that of a street brawl and a wrestling match. What makes those cases “borderline” is that the success state (or one of the possible success states) that are specified by a rule happen to coincide with what would count as success if the activity was *not* conceived as a game. The case of sonic pieces without “basic musical features” that are nevertheless counted as music in virtue of the background expectations that are standard for the appreciation of undisputed cases of musical pieces is different, in that what independently counts as “success” in creating a musical piece (specifying a sound pattern that exhibits tonal organization) is *not* achieved in those compositions. Yet, we can think of this in terms of the same conceptual distinction we made for the case of games. The kind of tradition where sonic pieces without “basic musical features” can count as music is a context where music-making can be thought in terms of the institutional theory of art, as a “game” with success states specified by arbitrary rules. Whereas sounds exhibiting tonal organization would count as music in *any* counterfactual circumstance (the experience of those objects involves the relevant NKA), not all contexts would allow for sonic pieces such as Ono’s counting as music. We must have a “rule” specifying that where a composer intends his work to be listened to against the relevant background expectation, *and* the work is accepted by the art/music world, then he has reached “success” in the “music game”. This context, though, depends on a distinction between music and non-music already being in place. The same point, given appropriate modifications, applies to art in general.

Finally, some counterexamples Kania raises against defining music as a subconcept of *art* (one of the “aural arts”) are purported cases of music which are not cases of art. Kania worries that Levinson’s definition excludes cases of “mundane music-making, such as the practicing of scales”,²⁴⁸ which arguably are not aimed at the “enriching or intensifying” of anyone’s experience; they’re not meant for an audience and have no aesthetic purpose; in sum, no characteristics that are central to their recognition as *art*. Anyone who has played Hanon or Czerny exercises at the piano for some time knows what kind of examples are meant here. (These are strikingly “unmusical” exercises, crafted to the sole purpose of developing digital skill.) Other examples are those of someone playing just to “warm up”, with no concern for the experience of an audience.

Kania holds that we should think of the concept of music and the concept of art independently, analogously to how we think of the concept of *depiction*: some cases of depiction are also cases of art, but many depictions are not artworks; some are not even intended as artworks but as, for instance, non-art illustrations (though many illustrations are artworks).

²⁴⁸ *Ibid.*, p. 10.

There are some strategies for the proponent of the definition of music as the art of tonal sound to deal with the force of such cases as that of the musician warming up or practicing scales, arpeggios, and the like. One of them is to bite the bullet and restrict the definition to those cases of music that *are* cases of art: where music is an art practice, it consists in the tonal organization of sounds (in the broad sense of “tonal” which includes purely rhythmic sounds); though there are cases of non-art tonal sounds that count as music. Now, the problem with this strategy for the way I’ve presented the definition of music is that I have classified music-making properties as a subclass of sonic aesthetic properties, *and* I have subscribed to a metaphysically more basic concept of art in line with aesthetic functionalism (a concept which co-exists and makes possible the institutional concept of art). And if the defining properties of music are already *aesthetic* properties, then the purported cases of non-art music are already cases of sound events “organized” to have certain aesthetic features, which is sufficient to make them art, according to aesthetic functionalism. So it seems that in accepting aesthetic functionalism *and* the idea that music-making properties are a subclass of sonic aesthetic properties, one is already committed to the view that there is no such thing as non-art music. Now, if there *are* really cases of non-art music, it seems to follow that one must either reject aesthetic functionalism or the idea that music-making properties are a subclass of sonic aesthetic properties.

Well, maybe not. As I’ve explained earlier, one of the consequences of endorsing aesthetic functionalism about art is that we get a category of art objects much more *vast* than the category drawn by the proceduralist concept of art. For many everyday aesthetic activities taking place outside the boundaries of the “artworld” (i.e. outside an institutionalized division of the arts) – cake decoration, quilting, tattooing, fireworks, etc. – will count as art according to aesthetic functionalism but not according to proceduralism. And the category of “artworld art” will also include things that are excluded by aesthetic functionalism (e.g., some, though not all cases of conceptual art). So, another strategy to deal with the non-art or “mundane” cases of music is to draw on the distinction between the aesthetic functionalist concept of art and the proceduralist concept of art. Under this distinction, those cases can count as art according to one concept but not according to the other, while the vast majority of cases will count as art under both concepts: a Beethoven piano sonata, for instance, satisfies both concepts. This would apparently preserve our initial definition from the counterexamples of non-art music. Though “mundane” music-making may not be “artworld art”, there is a wider category of “mundane art” (encompassed by the aesthetic-functional concept). But how plausible is the idea that practicing scales and arpeggios will result in something that counts as art, even with the qualification that it does so only under the “aesthetic-functional concept” of art?

Yet another strategy may be available, from our previous observation that the condition of organized sound events exhibiting basic musical features may be too weak – some things are to be heard for pitch properties or patterns of accented and unaccented sounds, while *not* counting as music, since those properties are not being used to generate *musical* relationships between sounds. (A pitch change may merely signal the illocutionary force of an utterance – we do not hear the pitch change as creating the sort of “virtual movement” of tones that Scruton specifies as the core of musical experience; we hear the pitch change as an “acoustic cue” for something that can be conveyed by some other vehicle in a non-aural medium, e.g., the graphic representation of a question mark.) Theodore Gracyk (2013: p. 9) has noted that “Although not every rhythmic, melodic pattern counts as music, rhythms and frequencies are the pre-cultural bedrock of musical experience.” By “rhythmic” and “melodic” patterns that fall short of being music he understands such things as birdsong, which we can *perceive as if it was* music because it exhibits the acoustic properties that sustain rhythm and melody, while not actually consisting of culturally mediated sonic designs, intentionally made not only to exhibit melodic and rhythmic patterns but also made in virtue of *insights* directed at the exploration of such patterns for aesthetic reasons, whether or not the aesthetic *result* is to be appreciated in itself (as in concert-hall music) or to serve some further function (as in ritual music). Birdsong is not song, nor music; no more than a cat’s meowing to express its content or discontent counts as music. The difference is that we can perceive birdsong *as if it was* music, while the same is not true of other animal sounds that do not exhibit rhythm-like and melody-like properties.

Applying this distinction to the purported cases of non-art music, we can see that one feature that counts in favor of their being *music* makes them differ relevantly from the case of birdsong: scales and arpeggios are not just non-musical sounds with rhythm-like and melody-like properties; they are the raw materials of music making. We organize the pitches comprised by the octave in twelve “equally tempered” tones not as “acoustic cues” replaceable for an equivalent cue in the same or another medium (whereas the natural function of birdsong might be equally satisfied by some other gesture of the animal, to signal a territorial claim or mating availability). In Gracyk’s words:

“No nightingale will ever select a Höfner bass in order to supply counterpoint to the song of another nightingale. Birds cannot choose their materials and thus their medium. (...) No bird will plan out ways to frustrate the learned expectations of the other birds in order to make a point to them”.²⁴⁹

²⁴⁹ Gracyk (2013: p. 34).

It is not enough to display the kind of acoustic properties that sustain “basic musical features” – a fact that allows for the distinction between the *music-like* and the genuinely *musical*. Partch was right against Cage: if you stop making *choices*, you stop making music; aleatory procedures may result in something that is *music-like* or that counts as music in an extended, proceduralist concept of music-making, but music in the primary sense must involve choices and insights about the aesthetic features of sound. To properly appreciate those works requires a background distinction between the musical, the music-like and the non-musical. Without such distinctions, aleatory sound events only count as music if by “music” we mean nothing but sound, as Cage apparently did, thereby emptying the concept, making it redundant. What distinguishes pieces like *Imaginary Landscape n°4* (for a set of twelve radios, operated by performers that handle volume buttons, according to the instructions) from truly mundane, non-art and non-musical sounds, is that we inevitably *frame* the decisions that go into the production of such pieces as making some form of “comment” on a pre-existent tradition of music-making.

Going back to scales and arpeggios. Take the example of Bach’s first *Prelude* from Book I of *The Well-Tempered Clavier*. It practically consists in nothing but a series of arpeggiated chords. If we take the piece apart and consider each chord in isolation, we have Kania’s examples of non-art music. What makes that sequence of arpeggios a case of art-music is the particular arrangement in which they were placed and the resulting aesthetic effect, the kind of musical movement generated by that organization. But in what sense are the isolated arpeggios already *music*? It can’t merely be in virtue of their status as music’s raw materials, since *any* sound can serve as music’s raw material. A rhythmic sequence built from the sounds of bouncing basketballs would count as an uncontentious case of music, but if we take the sequence apart and consider its elements in isolation, we no longer have musical sounds. The same is true of a series of pitches that constitute a melody – a single pitch sound will not count as music, except given an appropriate contextual setting (an ultra-minimalist “composition” of a single note may count as music in virtue of how it is to be experienced – what the appropriate “listening background” is). So what makes the isolated scales, chords and arpeggios *music*? My guess is that even taken in isolation, they already exhibit a kind of musical *movement*; they already afford the experience of sounds as *tones*, in Scruton’s terms.

Take a different example. If I start playing a C major scale on the piano and stop at B, I will frustrate your expectation of hearing the “conclusion” in the tonic (C an octave above the root). If I skip the repetition of the tonic and end the scale in D, a ninth above the root, you’ll get a feeling of

“open-endedness”, though not the “unsatisfaction” of ending in B.²⁵⁰ While not an impressive feat of aesthetic craft, I believe these basic operations minimally qualify for art in the aesthetic-functional sense. (On a related note: in one of his celebrated Norton Lectures of 1973, “The Unanswered Question”, Leonard Bernstein improvises a musical sequence based on one of Hanon’s scale-wise exercises for the piano.) If the two cases of frustrating the listener’s expectation have an aesthetic character then a normal sounding of the C major scale, ending in the tonic one octave above, equally has an aesthetic character, however little impressive – we hear a “virtual” movement of tones in it, and experience a feeling of conclusion or resolution in the end. Furthermore, the scale itself, as it unfolds, *suggests* possibilities of musical movement, in the way a chess move opens and closes alternatives. If La Monte Young’s *Composition 1960 #7* counts as art, which seems acceptable even for those who might resist counting it as music, I see no reason why a piece consisting of a single scale or arpeggio could not. Recall here Levinson’s example of the “colorful instantaneous combinations of chords of vanishingly brief duration”, which Kania himself is inclined to count as music, if they occur in the appropriate context of a music making tradition (and he does not entirely close the door to the hypothesis that a whole tradition of such sound practices could count as music, doubtful as it may be). I gather Kania would be inclined to count them as *art* even in those counterfactual situations where he wouldn’t count them as music.

Maybe we can use a simple distinction to shed some light on this matter. The distinction I’m thinking about is that between *types* and *tokens*. While the *type* “C major scale” or the *type* “C major triad on the root position” are *musical raw materials* but are *not* artworks, many *tokens* of those types may count as art, depending on contextual factors; just like the status of La Monte Young’s minimalist piece referred to above both as *art* and as *music* depends on such contextual factors. Another useful point is that of the *precedence* of art-music over the non-art components of musical “speech”. Consider a counterfactual situation where we had the same abilities we share with other animals, of perceiving sounds at the octave as equivalent, and other aural biases of key importance to musical perception. Imagine that we could identify pitched sounds, patterns of accented and unaccented beats, chords and scales, but there was no tradition of music making. In such a context, the “raw materials of music” would always be used for non-musical ends, such as the pitch changes signaling illocutionary force, or the patterns of duration in Morse Code. Arguably, in such circumstances, no scale, chord or arpeggio would count as music. What makes it intelligible for us to refer to scale exercises as “music” is that there is an art practice that takes scales and

²⁵⁰ I owe this example to sound artist Daniel Cioccoloni (see his interesting video, “Music and the Aural Landscape”).

chords as “ingredients”. Thus, Kania’s purported examples of non-art music have a *dependent* status on the practice of organizing sounds to exhibit those aesthetic properties based on the movement of “tones”. Otherwise, they could not be conceived as “musical” in any relevant sense. The music status of such things is as dependent on the *art* practice of music as the status of a move as a chess move is dependent on the game itself.

6 What musical works are

*Marco Polo describes a bridge, stone by stone.
“But which is the stone that supports the bridge?” Kublai Khan asks.
“The bridge is not supported by one stone or another,” Marco answers,
“but by the line of the arch that they form.”
Kublai Khan remains silent, reflecting.
Then he adds: “Why do you speak to me of the stones?
It is only the arch that matters to me.”
Polo answers: “Without stones there is no arch.”*
Italo Calvino²⁵¹

In the introduction to his 1989 book, *An Ontology of Art*, Gregory Currie nicely formulates the distinction between the “two distinct but complementary projects” of ontology and definition: “The first is to specify the kind of thing that art works are; the second is to specify the relation that things of that kind must bear to us in order for them to count as art works.”²⁵²

Throughout this work, I have been mostly concerned with the second of these issues, and I’ve tried to say something substantive about it, specifically as it applies to the case of music, by exploring the kinds of properties involved in the establishment of social and artifactual kinds, or in what it takes for something to be an example of such kinds. This was construed as a challenge to the Pragmatic Constraint (PC) on the metaphysics of art kinds. The challenge takes the form of a metaphysical constraint on the desiderata seemingly implied by the PC: these must themselves be evaluated in the light of metaphysical work on social and artifactual kinds, rather than being taken at face value as constraints on the metaphysics of art. This is not to deny that any constraints set by the PC are in fact legitimate constraints; but for those constraints that *are* legitimate constraints on art metaphysics it specifies that the reasons why they should constraint philosophical enquiry differ from the reasons typically subscribed to by defenders of the PC. That said, we should note that even the foremost proponent of the PC, David Davies, thinks of it normatively – we should think of artworks as things with properties *rightly* ascribed to them in the light of our critical and evaluative

²⁵¹ Calvino (2010: p. 74).

²⁵² Currie (1989: p. 2).

practices²⁵³ – which means that we are not to take at face value *everything* which is ascribed to artworks in practical and appreciative discourse, but that which *upon reflection* we deem to be rightly ascribed to them. In fact, such normative reading of the PC is what allows Davies to make a case for his own counter-intuitive ontological proposal, that artworks in general, and musical works in particular, are process-like rather than product-like entities, namely, action-tokens or what he terms “generative performances”, to be distinguished from the “focus of appreciation”, which is the work-product. So, that considerations from social ontology should be taken into account is not as radical a challenge to the PC as it might seem, when the PC is construed normatively rather than in a purely descriptive manner.

²⁵³ This point was underlined by Julian Dodd (2007).

6.1 The Puzzles of Musical Ontology

There is in the National Museum of Denmark an ancient Greek artifact, a grave stele in the shape of a pillar, dated from 200 b.C. to (most likely) c. 100 AD, originally discovered in 1883 by William Mitchell Ramsey, a Scottish archaeologist, near Tralleis, present day Aydin, in Turkey. It bears an inscription known as the *Seikilos Epitaph*. What is peculiar about this inscription is that it contains what is considered to be the oldest complete musical composition ever found. Along with the words that make up the poem, it has a form of musical notation, with pitches represented by letters and note durations by signs over the letters. The piece is an example of *epigram*, “a short and simple verse commonly placed on grave stones and votive tablets.”²⁵⁴ Here is one English translation of the Greek text:

As long as you live, shine

Let nothing grieve you beyond measure.

For your life is short,

*And time will claim its toll.*²⁵⁵

The song is attributed to a poet named Seikilos, of whom nothing is known besides what can be gathered from the artifact. It is presumed that the epitaph is dedicated to Seikilos’ wife, Euterpe, from what is inscribed after the poem: *SeikilosEuter* (“Seikilos to Euterpe”) though it could also be a patronymic (“Seikilos son of Euterpes”).

Though we have no way of knowing how ancient Greek music actually sounded like, musicologists have transcribed the piece to modern notation, and it has been performed and recorded. It is a small piece, approximately one and a half minute long. What the result of such “reconstructive” work allows us some kind of access to how the piece originally sounded like, when performed by musicians of the first century AD, is utterly obscure. If one plays the transcript tune

²⁵⁴ Mathiesen: (1999: p. 148).

²⁵⁵ The English translation, the original Greek text, a reproduction of the inscription with the ancient musical notation, as well as the transcription of the melody to modern musical notation can be found here: http://imslp.org/wiki/Seikilos_Epitaph_%28Anonymous%29

on a modern day keyboard, it will be rendered in a tuning system (equal temperament) different from what was used in Greek antiquity (just intonation). Nevertheless the melody as “reconstructed” is of a delicate beauty, with a haunting quality that is likely connected to the fact that we listen to it as an echo of antiquity; together with the fact that it sounds to us as in the (ecclesiastical) Dorian mode; and with the foreboding yet uplifting tone of the poem. Whether or not the transcription more or less faithfully renders the piece as it would be like in the first century AD, there is now a musical piece in its own right, to which we refer as the Song of Seikilos, a musical piece whose origins are connected with the physical artifact discovered in Turkey, 130 years ago, whether or not the musical piece as we know it (as it has been performed and recorded, since an early appearance, *ca* 1930, in the first of a set of twelve 78 rpm records, issued by Parlophone Records: *Two Thousand Years of Music*)²⁵⁶ can be said to be numerically identical with whatever musical piece is in fact represented in the ancient marble stele, and even if it turned out that the inscription was not the representation of a musical piece after all, and the efforts to decipher it were the result of utter confusion or wishful thinking. In other words, the Seikilos piece could be relevantly similar to Harry Partch’s *Two Studies in Ancient Greek Scales*, a modern work inspired in knowledge of antiquity. One feature that is distinctive of the Seikilos piece, in contrast with modern pieces evocative of Greek antiquity (by using the same scales, instrument types, etc.) is that the Seikilos piece as transcribed into modern notation is intended to be the presentation of a musical entity that actually existed. Though the Seikilos piece as we know it is causally connected with whatever musical entity is represented in the stele in Greek characters, it is not clear whether such causal link is identity preserving, so that numerically the same thing is represented by the ancient and modern notations.

I have described the Seikilos epitaph as a musical *composition* (our earliest specimen of such thing). Another way to describe it would be as a musical *work*. Is there a difference? That depends on how narrowly one construes the concept of *musical work*, or, to put it another way, it depends on how we *count* musical works. The concept of musical *composition* seems more neutral: it indicates that someone has made choices regarding the arrangement of sounds to be matched with the words of the sung poem, even if such choices are conceived more like rough guidelines for performance than like the strict prescriptions we associate with modern scores.

²⁵⁶ That recording can be listened to here: <http://gallica.bnf.fr/ark:/12148/bpt6k1275216>
 Other (more recent) recordings are: Petros Tabouris Ensemble, *Music of Greek Antiquity*, FM Records;
 Gregorio Paniagua, Atrium Musicae de Madrid, *Musique de la Grèce Antique*, Harmonia Mundi.

Some have argued (e.g. Goehr 1997) that applying the *work* concept to musical making prior to 1800 is anachronistic,²⁵⁷ in the sense that the concept did not attain regulative status until somewhere around that date, and so musicians of earlier periods would not think in terms of *works* – a kind of stable, enduring entity (with a set of normative properties) and clear conditions of individuation, unambiguously specifiable in a notation or score, faithfulness to which determines the *authenticity* of performances thereof. This would be a concept adequate to the rise to prominence of purely instrumental or “absolute” music, as autonomous from sung music or music with a text (autonomous in the sense that the music is not subordinated to a text)²⁵⁸. Others (S. Davies 2001) have suggested that we should have a more flexible concept of musical work, one allowing us to make sense of the continuities underlying the history of music, rather than focus narrowly on the idiosyncrasies of a particular period. In his book, *The World in Six Songs*, Daniel Levitin observes that “the right question to ask, in trying to understand music’s universality, is not what all musics have in common, but how they differ.”²⁵⁹ I find this an odd injunction. To understand how musical traditions (“musics”) differ, one must also understand what they share (to understand how musical traditions differ is to understand how they differ *qua* things of the same category – musical traditions). For instance, to understand how ancient Greek scales differ from medieval church modes; how these differ from major-minor diatonic scales, and how they all differ from scales used in other musical traditions around the world, we must understand what is common to all these cases: the concept of *scale* and what it refers to, its role in the creation and character of melodies, and many other features. Most music ever made, whatever its origin, is based on some sort of scale, the most common of which are pentatonic and heptatonic scales. To understand the wide-ranging differences within this phenomenon we must also understand its common core: the fact that melodies everywhere in the world are built from scales, from particular underlying organizations of tones that determine the character of melodies built from them. Generalizing, the understanding of musical *differences* is the understanding of different entities *qua* musical entities. It is one thing to understand how musical things differ from *non-musical* things; it is quite another to understand how musical things differ *between themselves*, i.e., as things of the same kind.

²⁵⁷ Goehr, (1989; 2007).

²⁵⁸ According to Lydia Goehr (1989: p. 55), the “[work] concept crystallized as a result of the fusion of four forces: (1) the articulation of the concepts of Fine Art and of the autonomous work of art in the mid to late eighteenth century and the subsequent inclusion of music under these concepts; (2) the emancipation of musical sound from poetry and the religious word, and the subsequent rise of absolute or purely instrumental music; (3) the specific and highly complex interplay between Enlightenment, Romantic and Idealist thought notably in German and French aesthetic theory; and (4) the emergence of a new sort of marketplace for musical works.”

²⁵⁹ Levitin, (2010: p. 6).

Following Stephen Davies on this point, I see no advantage in over-emphasizing difference and downplaying similarity. In light of that, I prefer to apply the same reasoning to the cross-temporal and cross-cultural application of the concept of musical work. To deny the existence of musical works in cultural contexts other than the “western” in the last 260 years seems like picking out a particularly “thick” conception of musical work, tied to a relatively recent historical context, to then conclude that the concept is not appropriately applicable to other contexts.²⁶⁰ In this respect, it resembles attempts to identify the concept of the aesthetic as an exclusively modern and “western” concept. During a debate concerning the cross-cultural applicability of the concept of the aesthetic²⁶¹, Jeremy Coote has aptly remarked: “Our opponents seem to want to choose the most ethnocentric definition of aesthetics they can find, so they can claim it cannot apply cross-culturally.”²⁶² This over-contextualization of the work concept is also in line with the propensity to treat the concepts of *art* and *artwork* as equally confined to the “modern” and the “western”.

Yet, are we forcing an alien cultural category onto something that it just doesn’t fit, when we describe the specific case of the Seikilos epitaph as a *musical work*? There is no straightforward answer to this question, I believe. To start with, the Seikilos piece is not one of “pure” or “absolute” music, but a hybrid artwork: it consists of words and a musical line (though in ancient Greece poetry was meant to be sung, and the Greek concept from which our concept MUSIC descends – *Mousike* – included poetry)²⁶³. To the extent that songs are musical works, the Seikilos piece is a musical work. But songs are ontologically more complex than purely instrumental pieces: they comprise at least a purely sonic or tonal structure and a sequence of words. We can always abstract the text from the music, even if the music is supposed to “emerge” from the text, not being originally conceived separately from the text that is sung.

When inquiring into music’s potential to convey, express, evoke or represent extra-musical realities such as emotional states, it is good methodology to focus on purely instrumental music, since in music with “lyrics” the representational properties of the sung text will most likely obscure

²⁶⁰ Davies (2001: p. 90) “[Goehr] over-contextualizes the musical work by taking features local to a particular musico-historical setting as essential to the work concept, thereby allowing the musical creations of other times and places to qualify as works only as a result of the imposition of a concept that is alien to them.” (*Ibid.*, p. 88) “The [work] concept is not historically recent. It is the conditions for satisfying it that altered and evolved, and these changes spread over many hundreds of years.”

²⁶¹ D. Dutton (1995) “Is Aesthetics a Cross-Cultural Category?”: http://denisdutton.com/weiner_review.htm

²⁶² Quoted in Dutton, *ibid.*

²⁶³ Hamilton (2007: p. 13): “Although ‘music’ comes from the Greek word *mousike*, classical Greece had no single word for what we call music; *mousike* (Latin *musica*) was used in a much broader sense than we use ‘music’. It literally means ‘the business of the Muses’, goddesses of poetic inspiration who sponsored every kind of intellectual or cultural activity.”

the waters, because we can never be sure whether a supposed example of expression, arousal or representation is not at least partly due to the words rather than the music alone. Similarly, when it comes to the problem of what musical works are, simplicity recommends that we should focus on purely instrumental pieces of music.

Though the Seikilos piece is not an example of such purely instrumental music, we still use “the Seikilos piece” or “the Seikilos Epitaph” to refer, e.g., a playing of the reconstructed melody without the words. We can certainly imagine arrangements of the melody on modern instruments, to which we would refer as “versions” of the “Seikilos piece”, a fact that constitutes one of the puzzling features of music.²⁶⁴

About the “Seikilos piece”, composer and documentarist Howard Goodall, in his recent book, *The Story of Music*,²⁶⁵ has made the following observation, which I reproduce here for the last three sentences and how they relate to the concept of musical work:

“The ancient Greeks, to be fair to them, did at least have a few tantalizingly scattered remains of a form of musical notation, the most complete example being some lines engraved on a first century burial tomb. This ‘Epitaph of Seikilos’ has an accompanying decipherable tune that lasts about ten seconds in all. But in general, it just would not have occurred to them that having musical notation mattered. This is because the Greek musical tradition, like all the musical traditions of the Ancient World, was one of improvisation. Setting music in stone, as it were, so that it stayed the same for each performance, year in and year out, would have struck them as a contradiction of music’s function and enjoyment.”²⁶⁶

Goodall’s observation about the predominance of improvised music in the ancient traditions and the pointlessness of “setting music in stone” is very close to the idea, also underlined by

²⁶⁴ Paul Ziff (1973: p. 722) has insightfully remarked: “The striking fact about the identity of pieces of music is I think this: a piece of music can have a variant but there is no such thing as a variant of a person or a stone or a pea in a pod. (At least not yet.) If you have a person and you have another person who’s sort of like him he’s a different person no matter how much like him. He’s a different person. You have two persons. And the same goes for stones and for peas in a pod. No matter how much alike they are if they are two then there are two of them distinct and individual. But if you have a piece of music and this other piece of music is almost exactly like it not exactly but almost exactly then it’s a variant of the first or the first is a variant of it or the two are variants of something else.”

²⁶⁵ A six episode TV documentary based on the book was aired on BBC Two.

²⁶⁶ Goodall (2013: p. 9).

Wolterstorff,²⁶⁷ of music-making in the absence of or before there are musical works (“music set in stone” one would say, perhaps with the qualification “more or less”.) This is perhaps what Lydia Goehr would describe as a tradition of music-making where the work concept is not regulative (which is compatible with the existence of something that we describe as musical works and perhaps coincides with Wolterstorff’s description of traditions where works of music exist not in virtue of someone having composed them, but where “the works just grew”).²⁶⁸

Yet, as the Seikilos piece perhaps testifies or at least prefigures, the interest in “setting music to stone” did emerge, and this in turn testifies to one of the most enduring features of the total human phenomenon which is music: an irresistible urge of “returning to the same”, so well captured in the expression “play it again”. This interest exists even where the “work concept” is not the object of a quasi-religious celebration that some authors point out in its deployment, from the mid-eighteenth to the twentieth centuries. Neither is such a reverential attitude a necessary condition for a work concept to be in place. (In Stephen Davies’ apt metaphor, the work concept admits of “thin” and “thick” variants.)

This “same again”, to which one returns in the listening experience, the *work* itself, is the subject of musical ontology, and the main question that it poses us is: what *kind of thing* is it? Another, no less important question is: are there really things of that kind? (The fact that we have a concept of musical work does not guarantee that such things *really* exist. A good analogy here is with the concept of *phlogiston* – we know what kind of thing it would be if it existed; it just so happens that it doesn’t.) But why are we interested in such questions? Isn’t the fact that we *experience* musical sound events as recurrences of “the same” enough? Well, for one, it is precisely this experience that is puzzling: what we experience in multiple performances of a musical work are spatiotemporally distinct *events*, more precisely, *sound events*. Each of these sound events is a particular, concrete, unrepeatable entity. It comes into existence at a given moment in time and it ceases to exist at a later moment. Recordings make the picture more complex. When we listen to multiple playbacks of a particular recording of a particular performance, what we experience are also spatiotemporally

²⁶⁷ (1987: p. 115) “Notice to begin, that the presence of music in a society requires neither composition nor works. Before ever works of music were composed, there were works of music. Though people in traditional societies have works of music, mainly they do not have them by virtue of someone having composed them. The works just grew. And more basic even than this is the fact that before ever there were works of music, there was music.”

²⁶⁸ Notice that the concept of “composition” deployed by Wolterstorff in the passage above is thicker than the one I used, when I remarked that it seems “more neutral” than that of musical work. The difference may be cashed out in terms of there being someone with a special social role – the composer – and the thinner notion of choices having been made, individual or collective, regarding the organization of some “musical material”.

distinct events, though each playback reproduces the *same* performance. Yet, there is a difference between experiencing the performance, which requires presence at the original event, and experiencing playbacks of recordings which encode the sounds of the performance. Playbacks are not recurrences of the original sound event, whose acoustic features are encoded in the recording; rather playbacks are distinct sound events which are also (better or worse) acoustic clones of the original event. (Further complications are introduced by sound manipulation in recordings, in virtue of which a recording may not correspond to any actual performance that has taken place in the usual way.) For simplicity's sake, let us set recordings aside and stick with performances – actions of playing musical instruments according to instructions specified by scores, or emulating a “paradigm” performance (for traditions where a musical notation is not available).

So, there is this common element in the experience of music, whether that experience is one of music-making seemingly without works (e.g. imagine a tradition of purely spontaneous improvisations)²⁶⁹ or one where there is an interest in re-encountering “the same” (the work) in multiple performances, namely: those experiences are always experiences of unique, unrepeatable sound events. One important distinction is that in the latter scenario some of those sound events are experienced as *recurrences* (or tokens, or instances) of one single entity – a musical work, which we pick out with a name or referring expression, such as “the Seikilos piece”. This allows (in fact, it requires) us to notice a subtlety in the concept of *performance*: both spontaneous improvisations and recurrences (tokenings, instantiations) of a musical work are *performative actions*, but only the latter are *performances* in the sense of being performances of something distinct from the performative action itself – the work that is *performed*, i.e., rendered or presented in performance.

We have thus arrived at what is supposedly the main puzzling ontological feature of musical works: their *repeatability*; the fact (if it is a fact) that we encounter one and the same work in multiple performances, and that this is the only way to fully experience a musical work: by having aural experiences of one or more of its legitimate instances (its performances). Of course, by describing musical works as repeatable we are already making an ontological assumption: that there are repeatable entities. So, what first puzzles us about musical works is not so much an established ontological *fact*, but rather the fact that the most natural response to the experience of some

²⁶⁹ Supposing that a purely spontaneous improvisation does not constitute a work. Opinions are divided here, between those who think that the fact that an improvisation *could* be preserved, in notation or otherwise, is enough for that improvisation to constitute a work, and those who think that it is not sufficient. If the latter are right, then (at least some) improvisations are a distinct kind of musical entity from works.

performed sound events is to posit the existence of a certain kind of repeatable entity, i.e., the musical work.

Maybe this is not the right way to describe what really goes on. It's not that we have some aural experiences and *then* posit the existence of a kind of repeatable entity – the work – as if to explain certain features of the experiences (e.g. that we experience those sounds as the *recurring* of something); rather it seems that the feature of repeatability is already built-in the concept of a musical work, and presupposed in the very distinction between a work and performances *of* a work. (Whether the concept, with that built-in feature, actually picks out something real, is a distinct issue.) Now the source of our puzzlement takes the form of a conditional: *if* there are musical works, they are repeatable entities. Repeatable entities, if there are any, are ontologically puzzling (repeatability is an ontologically puzzling feature).²⁷⁰ At this point, the problems of musical ontology begin to appear more clearly: 1) we must account for our experiences of “the same again” that ground our use of the work concept; 2) this will be done either by (a) vindicating the work concept as the concept of a repeatable entity, in which case the main problem of musical ontology will be to explain musical works' repeatability (and such explanation will vary according to what *kind* of repeatable entity musical works are identified with), (b) reconstruing musical works as some kind of non-repeatable entity (e.g. a concrete particular, a set of *concreta*, a fusion of *concreta*) in which case repeatability is to be explained *away* rather than explained; (c) eliminating musical works in favor of an ontology that appeals to performances only (this option may involve an error theory concerning our statements about musical works).

To get a feel of what is puzzling about the ontological status of musical works one need only compare two situations where the same question, about spatial location, is posed about two different sorts of object. To the question “Where are the Elgin Marbles?” the answer is straightforward: “In the British Museum.” In fact, the straightforwardness of this answer is what makes intuitive the identification of non-cast sculptures and paintings with concrete physical objects. To the question “Where is the Seikilos piece?” the answer is anything but straightforward: is it in the National Museum of Denmark (where the marble stele with the original inscription is located)? Is it wherever a notation of the “reconstructed” melody is? Is it wherever a performance of the melody occurs? Is it in the minds of people who experience the performances or read the notation? Or is the question

²⁷⁰ I'm not suggesting that our puzzlement about repeatable entities is dependent on the existence of repeatable entities. Part of what makes repeatability an ontologically puzzling feature (whether or not there are things with that feature), is aptly expressed in the question “How can there be such things?”

simply malformed, resting on a category mistake: musical works are just not the kind of things that exist in space? In which case the answer to our question would be “Nowhere!” But if musical works exist nowhere in space, and yet they do exist, what kind of things are they? And if they exist wherever their performances exist, it seems that they can exist simultaneously in different places, not in the sense that an object exists in more than one place by having spatial parts that occupy different regions of space; nor in the sense that a scattered object occupies non-contiguous regions of space, by having disjointed spatial parts, perhaps at a remote distance from each other; but in the sense of existing *wholly* at distinct spatiotemporal locations. This seems like a puzzling consequence of the fact that there can be several performances *of the same work* at different locations, together with the idea that works are located where their performances are: that one and the same thing – the work – can be located at a certain distance from itself.

Other strange consequences emerge from the fact that performances of works exist at different *times*: if works existed *only* where and when their performances exist, then works would exist intermittently; they would constantly pop in and out of existence. Under this conception, works would not come into existence in the act of their composition – a composer would produce a score of a work but its existence would only begin at the moment of the first performance. But in that case, what entity would the musical score represent? What entity would be the object of the composer’s imaginative and compositional actions? Besides, to perform the work for the first time, the performers would have to have some epistemic access to the work *before* they engaged in the actual bodily actions constitutive of the performance. And what about the status of works that are never performed (though they have been composed or indicated via a score)? Should we say that they don’t exist? Do the different scores of unperformed works represent the same thing, namely, nothing?

These awkward questions set before us the problem of what are musical works to be identified with, if they exist at all. What are these things we call musical works? Among the several candidates to fit that role are the following: the original score; the set of all scores; the first performance said to be “of” the work; the set of all performances that accord with the score; the mereological fusion of all performances; a mental item or idea in the composer’s head; a mind-independent abstract entity to which scores, performances and mental representations or ideas bear some appropriate relation; the compositional action through which the composer arrives at the “final product” (the score or a first performance). Each of these hypotheses has its own difficulties answering the perplexities that arise when we ask seemingly simple questions such as “Where is the Seikilos piece?”, and each generates perplexities of its own. (As it often happens in philosophy, we

appear to solve one problem, only to give rise to another, and there is no straightforward way to eliminate all of them or to choose between rival theories.)

This problem of what musical works are to be identified with is also known as the “categorical question” (Dodd 2007: p. 1) – the question of what *ontological category* musical works belong to. So, for instance, if we were to identify musical works with their scores, we would (so it seems) be ascribing musical works to the category of *concrete particulars*: things that have spatiotemporal existence, of which medium-sized physical objects are the most likely examples to come to mind; though even here the problem arises again in another form, since we now must ask in virtue of what every physical token of a work’s score is a token of that and not another score, i.e., scores too seem to be “repeatable” entities, no less than musical works. If, on the other hand, we were to identify musical works with performances, we would be ascribing them to a subclass of concrete particulars called *events*: things that are, so to speak, *extended in time* and as such have *temporal parts* (e.g. the 15th minute of a musical performance is a temporal part of that event – the performance). But to complicate matters, there are at least two potential ambiguities in the idea of identifying works with performances. First, there is an action/product ambiguity in “performance”, by which we may either be referring to the performative actions of musicians (actions being a subclass of events, with the added problem that it is not entirely clear what distinguishes them from events in general) or to the product or result of such actions – the sounds that issue from the instruments and together constitute the sound event which we call “performance”. Typically, when we speak of performances of musical works, what is meant is the sound event produced by the performative actions of musicians; though sometimes the boundary is fuzzy.

Another ambiguity in the idea of identifying works with performances is a type/token ambiguity. Each individual performance is a concrete particular event, a non-repeatable entity which is extended in time. But when several distinct events conform to a pattern, we also speak of *event-types*: for instance, that which is repeated in what we call “repeated actions”. My action of flipping a coin at t_1 is an *example*, *instance* or *token* of the event-type *flipping a coin*, of which your action of flipping a coin at t_2 is another token. But while a token-event is clearly an *event* – a concrete particular extended in time – it doesn’t appear to make sense to think of the event-type itself as a concrete particular extended in time. The event-type is more intuitively identified with the pattern to which all token-events of that type conform, and in virtue of which they are said to be *of* that type;

much in the way that, for instance, observing two identical (same-shaped) bridges we identify a certain geometrical structure or pattern they both exemplify or token.²⁷¹

One hypothesis concerning the nature of musical works is that they are *abstract entities*: mind-independent things that do not exist in space but have spatiotemporal “instances”, “occurrences” or “exemplifications”.²⁷² The parallel here is with the familiar category of *universals* (properties and relations): take for instance the property of redness. It can be shared by several concrete particulars, so that a realist about universals will either say that redness is located nowhere, though it can be *instantiated* or *exemplified* by objects in multiple spatial locations, or that it is (wholly) located where its instances are located.

A key difference between universals (conceived as mind-independent abstract entities) and musical works (likewise conceived) is that the former, but not the latter, are *predicated* of particular things (including other universals): we say that the table cloth is red (it has the attribute of redness) and that red is a color, shareable by multiple individual things. Yet, the property of redness, or being red, is not itself red. Red things (such as the table cloth) have the property of redness, in virtue of which the predicate *is red* applies to them, but those concrete individual things are not themselves predicated or attributed to other individual things. In other words, the object denoted by “table” differs crucially in this aspect from the property *is a table*, which is shareable by many distinct individual things denoted by “table”. Musical works seem more akin to particulars than to properties: things that are the bearers of properties (or subjects of predication) but are not themselves predicated of other particulars, in the way redness is predicated of many distinct objects.

The difficulties involved in conceiving works of art taking universals as our model for understanding were summarized by Eddy Zemach in an article on the ontology of interpretation.

²⁷¹ As expected, even here things are not as straightforward as they might seem. For instance, Eddy Zemach (1970; 1982; 1997: pp. 140-166) defends a conception of types according to which they are conceived as concrete entities. He also defends that all existents are types, i.e., things that *recur* in time or in both space and time. In his conception, for instance, both Socrates and The Greek are types; and they both are (contingently) identical at an *index*, in this case, Socrates.

²⁷² There is a subtlety regarding whether abstract entities are to be conceived as timeless existents (existing at no time) or as eternal existents (existing at all times). Julian Dodd (2007) argues for the latter position on the basis of something like the following reasoning: if a musical work can have performances at time *t*, or be thought about by me at time *t*, then it must exist at *t*. The antecedent is obviously true, so the consequent is true by *modus ponens*, which defeats the conception of musical works as timeless abstract entities.

About the identifying works of art with universals or any ontological constructions deeply resembling universals in their key aspects, Zemach says:

“[T]hat solution is entirely unacceptable: it is obvious that artworks are not universals. Universals are not material things. They do not occupy space and do not take time; they have no mass or weight, and they cannot be physically handled. No universal can be damaged and repaired, be created and destroyed, be heard, seen, or touched. Yet these things are true of works of art.”²⁷³

Nonetheless, this solution of identifying at least some works of art with universals or abstract objects of some kind has appealed to many as the solution which, despite having its problems, it may be the one which generates the least problems, while solving quite a few perplexities. Of course, not everyone will agree to that description, depending on how strongly counter-intuitive the “revisionism” of our pre-theoretical beliefs imposed by Platonism will seem to one theorist or other.

The theory known as Musical Platonism, in its most currently developed form, was defended by authors such as Peter Kivy and Julian Dodd, based on contributions by Richard Wollheim and Nicholas Wolterstorff. It is a theory with an immense intuitive appeal, since it seems to explain in a straightforward way the feature of repeatability of musical works.

Platonism conceives musical works on the model of putative abstract entities such as properties or universals – the standard “repeatable entities” of general metaphysics – yet differing from properties in the sense of being particular individuals, not predicable of other particulars; and differing from standard particular individual things – which are spatiotemporal and concrete – in the sense of being abstract and non-spatiotemporally situated. In other words, Platonism views musical works as eternally existent (existing at all times), abstract (without spatial location or spatial and temporal parts), non-predicable, individual things. The name standardly adopted for the ontological category comprised of such entities is “type”, following a well-known distinction between types and tokens, of which the standard examples are those of word and letter types and their tokens, i.e., between a given word or letter and the concrete manifestations or occurrences of those words and letters. For instance, in the word “discombobulated” there are two token-letters or occurrences of the letter-type “o”, and if I write that word five times on a page, there will be five word-tokens but

²⁷³ Zemach (1990: p. 297).

only one word-type, namely, five tokens of the word-type “discombobulated”. The type-token distinction fits nicely with the notion of a musical work which is “manifested” or exemplified in its performances, which are the tokens of the type. In a Platonist reading, performances of musical works are the concrete tokens of the abstract work-type, the latter being understood as a type of sound event, or, a performed sound event, if one wishes to exclude the bizarre but nevertheless real possibility of random (or non-intentionally produced) sound events which accidentally exemplify the sequence of sounds specified by the type from counting as bona fide instances of the work-type.

According to Musical Platonism, thus, a musical work is identical with a “sound structure” or a type-sequence of musical sounds or tones. Another way to unpack the notion of sound structure is to describe it as a rhythmic, melodic or harmonic pattern. At this point, Platonists divide between those who adopt the view that the sound structure to be identified with musical works includes timbral properties besides the rhythmic and tonal relationships which together constitute the relevant sound structure (e.g. Dodd: 2007). This view, which takes timbral properties as equally normative for musical works is known as timbral sonicism, while those who do not count timbral properties as normative (e.g. Kivy: 1993) are known as pure sonicists. Sonicism, in both its pure and timbral varieties, is a thesis about the individuation of musical works, i.e., a thesis about which properties are normative for the production of bona fide instances of the work-type. In this view, two concrete sound events, *S* and *S**, can only be counted as bona fide instances of the same musical work *W* if they have those properties which are normative within the type. This is compatible with a degree of vagueness in the realization of such properties. For instance, that a certain work should be performed in a given tempo, say, *adagio*, may be normative for that work, but the property itself is vague, so that different renditions of the piece in *adagio* may nevertheless differ slightly in length and momentum. The same applies to timbral properties, since instruments of the same type, producing token sounds of the same timbral type (e.g. two oboes) will nevertheless exhibit differences in tone color.

It is important to distinguish timbral sonicism – the thesis that timbral properties are normative within a sound structure identified with a musical work – from a quite different thesis: instrumentalism – the thesis that the instrumental means of sound production are normative for a given musical work. According to instrumentalism, if a work *W* was written for the oboe, it does not suffice, to produce a bona fide instance of that work, that a sound event intended as an occurrence or performance of that work incorporates sounds which are oboe-like (but, instead, produced using a Perfect Timbral Synthesizer); for a defender of instrumentalism, a bona fide token of such work would have to be produced using an actual oboe. Instrumentalism is opposed to

sonicism, in that it takes non-acoustic properties as normative within musical works. If the work is to be identified with a sound structure, then only acoustic properties can be taken as normative. To accept non-acoustic properties as normative for musical works is to depart from sonicism, though not necessarily from Musical Platonism (a qualified version of Platonism, endorsed by Levinson, incorporates instrumentalism), but any form of Platonism which accepts instrumentalism will not identify the work tout court with a sound structure.

The term “sound structure” may be misleading since it is ambiguous between two readings, one of which suggests that the type itself is structured in the same sense that the tokens of the type are structured, i.e., in the sense that they have spatial and temporal parts. Dodd (2007) has convincingly argued against such an interpretation of “sound structure” when used to describe work-types qua abstract entities. This is because abstract entities are not the kind of thing which can have spatial or temporal parts, since they are neither concrete particular events nor are they located in space, so that they could be made up of spatial parts bearing among themselves the same kind of relations that concrete particulars exhibit among their spatial parts. As such, the attribution of a “structure” to types qua abstract objects should be taken in the second reading, one which sees at least certain predicates attributed to types according to a doctrine of “analogical predication”, i.e., the thesis that there is a systematic relationship between the predicates attributed to a type and the predicates attributed to the tokens of that type, but the properties expressed by those predicates are not the same. So, for instance, while the predicate “has a C sharp in its seventh measure” is applied literally to tokens of a work-type, the type itself cannot literally have that property, since, as an abstract entity, it has no “structure” in a literal sense, i.e., it has no temporal or spatial parts. Rather, the predicate as applied to the type expresses a different property, namely, the property *is such that any bona fide instance of it has a C sharp in its seventh measure*. In other words, only performances or bona fide tokens of the work-type instantiate literally properties like *has a C sharp in its seventh measure*. The type specifies the properties which are normative for any bona fide instance, but not all predicates which properly apply to tokens express the same property when applied to the type itself. The doctrine of analogical predication can thus be used by the musical Platonist (in fact, it seems that it *must* be used) to explain another of the crucial features of musical works, namely, their audibility. This is because one objection that might be raised against construing musical works as abstract entities is that abstract entities, not having temporal or spatial parts, cannot be directly heard, in the way that performed sound events can. So when we say that a musical work, qua abstract type, is audible, we are saying that the type is such that its tokens are audible entities. Another way to put this is to speak of “deferred ostension”: we hear a musical work *by* hearing one of its tokens.

Despite its intuitive appeal in the way it purportedly explains the repeatability of musical works, Platonism clashes with a naturalist sensibility suspicious of the category of abstract entities or any putative existent which lacks spatiotemporal location. If we share such naturalist sensibilities, we are prone to ask the question “But if such things as abstract types exist *where* are they?” Of course, the question is misplaced, since if abstract objects are defined, by contrast with concrete particulars, as lacking spatial location, it makes little sense to ask where they are. But this very fact is what makes them suspicious: how can entities which do not exist in space have any kind of *effect* on the experienceable world? Musical works thus conceived disconcertingly resemble such exotic beasts as Platonic Ideas (which accounts for the doctrine’s name) inhabiting a Platonic Kingdom, causally disconnected from the physical world. Yet ontological parsimony seems to advise us against postulating such entities if we can avoid it without theoretical loss. Philosophers like Dodd try to make the Platonic picture more amenable to a naturalist sensibility by invoking a conceptual dependence between abstract types and their concrete tokens: according to that view, types are *intrinsically* of their tokens in the way that chess moves are intrinsically of chess pieces, though no chess move is existentially dependent on any particular concrete chess piece. Likewise, musical works qua abstract types would be intrinsically of performed sound event tokens, though no musical work is existentially dependent on any particular concrete performance of it. It seems utterly plausible that no musical work depends for its existence on having just the particular performances it has actually had since its composition.

Amenable as this makes the category of abstract types, making us think less of them as denizens of a Platonic Kingdom, wholly disconnected from the natural realm, it does not quench the suspicions it raises in the naturalist minded. Whether or not conceptually tied to their concrete tokens, abstract types are still ontologically bizarre or perplexing creatures, even though they fall under an established ontological category, one with a long-standing history in metaphysics.

However, the problem of making abstract objects amenable to a naturalist worldview, which would be partly achieved, for instance, by showing they can somehow participate in causal relationships, or that their postulation is required for an intelligible picture of the world, is not a problem to be solved by the special metaphysics of musical works or of artworks, but rather a problem for the general metaphysician. The focus for the ontologist of music, instead, is to determine whether we have reasons, from within our discourse and understanding of musical works, to subsume them in one ontological category rather than other, whether there are reasons to postulate a *sui generis* category for musical works, or whether we should dispense altogether with

the idea that musical works have any place in our ontology, that is, whether musical discourse requires, for its usefulness, that musical works are thought of as part of the furniture of the world.

My aim here is rather more modest. I want to look at some reasons why musical works should not be identified tout court with abstract types, or “sound structures” conceived as a kind of abstract entity. This is, to a certain degree, independent from the question of whether the world contains such things as abstract objects, and whether or not abstract objects must figure in a description of the fundamental nature of musical works. This is not to say that it isn’t interesting for the musical ontologist to determine if talk of abstract objects should be absent from the description of the nature of musical works. After all, eliminating abstract entities from our picture of the world would substantially change our picture of musical works, namely, whether or not the description of their nature would include sound structures qua abstract objects. Yet, there is a difference between the issue of sound structures or sonic patterns being best conceived as abstract entities, and the issue of work-ontology providing us with reasons not to identify musical works with sound structures or sonic patterns tout court.

The most counterintuitive consequence of identifying musical works with abstract sound patterns is that this seems to place musical works outside the scope of human creation. While it is intuitive to say that Pythagoras discovered the theorem that bears his name, not having created the proposition expressed by any linguistic formulation of it – this was already true before his discovery, and could very well have been discovered by someone else – and it is equally intuitive to say that no geometrical figure existentially depends on any concrete object exemplifying that figure, it seems to go against the grain of our intuitions to say that composers discover, rather than create, the musical works they compose and that another composer, working even in a different musico-historical context, would compose numerically the same work if he produced a score specifying the same sequence of sounds or tones. The counterintuitiveness of this idea was explored in Borges' tale of Pierre Menard, as we have previously seen, and it finds some resonance in Danto's thought experiments of structurally identical objects which nevertheless differ significantly in their contextual properties, to the point of that making it the case that we have numerically distinct works of art and even cases of structurally identical objects which differ in art status (some of them being artworks and others not). If the intuitions behind such thought experiments are true, then context plays a role in the individuation of both artworks generally and musical works in particular.

This conflict between our pre-theoretical intuitions and the postulation of art abstracta rises with special force when we consider the “repeatable arts”, like music, that is, arts or art forms whose objects – the artworks – have multiple “instances” and are not identifiable with any given concrete

particular object. For instance, if two sculptors working independently of one another, in different times, produce non-cast sculptures which are structurally identical, our intuitions tell us that they have produced distinct artworks, and not different instances of numerically the same artwork. And the conflict does not arise because there is no pressure to identify non-cast sculptures with an abstract structure they exemplify. We do not identify Michelangelo's *David* with the structure instantiated by the marble object standing in front of the Palazzo della Signoria in Florence, hewn by the artist between 1501 and 1504. Rather, it is the marble object itself, that specific marble object, we identify with the artwork created by the artist. No other structurally identical marble object would be considered by us as numerically identical with the original artwork.

On the one hand, the kind of intuitions at play that tie in our minds the concrete object, the artwork and the particular creative actions of the artist, also operate in the case of multiple (or "repeatable") artworks such as music. We tend to think of Beethoven's *Fifth Symphony* as tied to the particular creative acts of its author, as something we do not properly understand as an artwork unless we understand it as issuing from the particular creative or artistic context where it did rise or became epistemically available to us. We see it as occupying a certain place in the composer's *oeuvre* and bearing significant properties in virtue of occupying that place. We also see it as occupying a certain place in the history of music and bearing significant properties in virtue of occupying that place. If we hear a structurally identical piece and learn that it has been arrived at by a different composer, working in a different musico-historical context, it would seem to us to have some very different artistic and perhaps aesthetic properties, similarly to what happens with the fictional Pierre Menard text. On the other hand, the fact that we don't identify the musical work with any particular concrete object gives rise to a conflict between these intuitions and the equally intuitive identification of the work with the sound structure instantiated by performances of the work. Though intuitive, this identification, if true, has consequences that imply the negation of the former intuitive notions: that the musical work is a creation of its composer, not existing previously to the act of composition; that it is tied to particular creative acts, within a particular creative or artistic context, and bearing historical relations to other artworks and non-artistic events that are relevant for its artistic and aesthetic profile and thus to its identity.

It is as much a problem for musical ontology to solve this clash of intuitions as it is to find an ontological category that plausibly fits musical works. We have thus the problem of deciding which intuitions should have any metaphysical weight, or whether some of them may be sacrificed with a minimum of theoretical loss. Philosophers friendly to a pure Platonist approach (as Kivy and Dodd), intent on preserving its theoretical virtues and the apparent simplicity of its explanation of

musical works' repeatability, downplay the weight of our intuitions concerning the creatability and contextual individuation of musical works. According to them, to think of musical works as discovered rather than created by their composers is made amenable to our sensibilities and intuitions by the thought that such acts of discovery are nevertheless creative acts, not in the sense that they bring something into existence, but in the sense that performing those acts requires creativity from the part of composers who, in that manner, creatively discover the sound structures that the Platonist identifies with musical works. However, this "solution" still leaves us with the incredible result that the *Fifth Symphony* existed for many centuries before its "creative discover" by Beethoven. Such idea clashes with our intuitions even if it seems plausible to suppose that the purely structural relations of sounds, specified by the work's score, do not existentially depend on the composer's creative or compositional actions. They were somehow already present in the "logical sonic space" and the idea of the composer "discovering" such relations in a way similar to how the sculptor "finds" certain shapes in the marble is not an incredibly bizarre one. Yet, there is something bizarre in the claim that the musical work pre-existed the acts of composition and indeed already existed even before there were human beings. Even if we think of the artist's actions as an instance of rearranging previously existent matter (as in the case of hewn statues), such that strictly speaking the world is not fundamentally different from how it was before, we do in fact resist the idea that either the creative actions of composers or the creative actions of sculptors (and other producers of "singular artworks") do not introduce anything new into the world. There is a sense in which the world is in fact different from what it was before the artist brings a work into completion, making it the case that henceforth something previously unavailable – the artwork – is now available to us. The social, historical and artistic landscape is indeed modified by the fact that such works are made available.

Such considerations have lead ontologists of music to pursue different solutions to the problem of the nature of musical works. Some of those solutions seek to preserve the theoretical advantages of the type/token distinction, without identifying musical works tout court with abstract sound structures. For instance, Jerrold Levinson (1981; 1990) proposes an ontology of musical works according to which they are the contextualized result of actions of "indication" by their composers, that is, the selection and specification of a structure, from the infinitely many available in "logical sonic space", such that the musical work is not identical with the sound structure per se but with the contextualized product of that act of indication, or, in other words, with the sound structure as indicated by that composer in a given time or musico-historical context. This proposal has given rise to many objections, in particular to the objection that accepting such entities in our ontology would lead to an undesirable proliferation of "indicated structures" (Levinson's technical term) or,

to put it in a general form, “X as indicated/discovered/selected by P”, where P stands for a person or epistemic agent. For instance, according to the objection, “penicillin as discovered by Fleming” would be an entity distinct from penicillin independently of Fleming’s act of discovery. This is a serious objection, though, as we will see in the next section, there are responses available to it.

It is not my purpose to evaluate each and every ontological proposal made to this day, in the effort of solving the “categorical problem” of the nature of musical works. These have been abundant and gave rise to many interesting discussions, since, as befits a philosophical theory, each proposal seems to solve one problem, or one particular aspect of the global problem, only to raise further difficulties at some other point.

Some authors, like Gregory Currie (1989) and David Davies (2004) have made ontological proposals that identify musical works not with sound structures but with the compositional actions that take such structures as their object. Such proposals go against our intuitions that the work is rather the product of the creative actions and not the actions themselves. In the specific case of musical works, it is argued that actions are not the appropriate category to subsume musical works, since these have properties that musical works, as commonly conceived, could not possibly have (e.g. a compositional activity may be interrupted but no musical work has such property) and musical works have properties that don’t seem to be shared by compositional actions (e.g. being audible or having a C sharp in its seventh measure). Similar objections have been raised against proposals that identify musical works with mental entities of some sort, such as ideas in the minds of composers.

In what follows, I’ll focus on the ontological proposal that identifies musical works with some sort of structure, the main division being between those who identify musical works with pure or abstract sound structures and those who identify musical works with some kind of contextualized sound structure. In particular, I’ll focus on Robert Howell’s elaboration of a contextualized ontology for musical works, based on Levinson’s theory of “indicated structures” and “initiated types”. I will try to show how combining Howell’s distinction between *properties*, *patterns* and *types*, Levinson’s notion of “indication” and Searle’s apparatus of “status-functions imposition” can be conceptually illuminating of our art practices and what an ontological mapping of such domain must be like.

6.2 Musical Works and Social Reality

Layers of intentionality-dependence

By the “ontology” of an item I mean the junction of two separate aspects: a) the “ontological category” to which that item belongs or, in other words, what *kind* of thing it is; b) that item’s conditions of persistence and individuation. These are the properties in virtue of which that item is *numerically one and the same* across times or in virtue of which any two particular objects or events count as *tokens* of numerically one and the same item. An item’s *essential properties* are a subclass of its persistence and individuation conditions, that is, those properties the item has in every possible world in which it exists but that stop short of *individuating* it. By the notion of an item’s ontology “making reference” to any kind of property I mean *either* that the ontological category that item belongs to is (at least partially) determined by the item’s possession of that property *or* that the property in question must figure in a specification of the item’s conditions of persistence and individuation.

We have seen that part of what makes musical works intriguing as a subject of ontological enquiry is that (whatever else they are) they are intentionality-dependent, i.e., their existence depends, in one way or other, upon the *directed* mental states of human or other agents. Intentionality-dependence adds a layer of complexity to our understanding of the nature of any entity having that feature. This is true even if we arrive at the conclusion that intentionality-dependence is no more mysterious than dependence on any non-mental phenomenon. The ontology of intentionality-dependent entities is *prima facie* less clear to us solely and precisely in virtue of their having that property (being intentionality-dependent), regardless of whether that property turns out not to be mysterious in itself, or adding any mysterious component to the ontology of those things that exemplify the property. The initial obscurity is thus epistemic and not metaphysical, in any obvious way.

Having said that, we should now provide a brief sketch of the way or ways the property of intentionality-dependence obscures the ontology of one or other kind of entity. A few questions immediately spring to mind: in how many ways does intentionality-dependence obscure the ontology of items exemplifying it? Or is it always obscured in the same way, or in virtue of the same feature or set of features supervening (at least) on that very property? If so, what features are these?

We have studied, in previous sections, the nature of social and institutional entities (such as games), institutional entities being a subclass of social entities. Both social and institutional entities

are in their turn a subclass of intentionality-dependent items. This suggests that there are at least three great layers of complexity in the phenomenon of intentionality-dependence, such that the “upper” layers subsume the lower ones:

1) The level of “raw” intentionality-dependence, i.e., things that are dependent upon the existence of mental states directed towards them but do not involve any kind of joint action or shared representations of how things stand in the world. For instance, beliefs are intentionality-dependent items, for nothing is a belief unless there is a cognitive agent in some state which is “directed” towards some object; yet, the ontology of beliefs, whatever else they may be, does not involve any reference to joint action or shared representations.

2) The level of social or socially constructed (but not institutional) things and events. Here we find those intentionality-dependent items whose ontology must make reference to some form of joint action or shared representations but have no functions that depend crucially on a shared representation of the function itself or the item’s *status* (meaning: the function does not require conceptualization in order to be the function it is and to be carried out). One example of such is that of two people going for a walk or singing a duet: the action is social in that it is not merely the summation of individual movements but involves each participant representing his own actions as *part* of a collective action (*we do X*): that of taking a walk or singing a duet. But neither the walk nor the duet *count* as a walk or as a duet in virtue of any shared representations about walks or duets. What is shared here is the “*we do X*” form of representing one’s actions.

3) Finally, the level of socially constructed things and events which also have institutional status. One example of such is language; given the correctness of the view that language is essentially non-private, its ontology makes reference to the joint action and shared beliefs of human (or other) agents. Given that the functional properties of words and sentences supervene on more than the perceivable intrinsic properties of any actual fragments of speech, and that this something more includes the mutual recognition of function by a community of speakers, language belongs to that class of things which are essentially socially constructed and institutional. The ontology of language makes reference to joint action and shared representations of cognitive agents, as well as to status-functions. (In fact, if the approach I have taken throughout this work is correct, reference to status-functions in the basic ontology of a thing is the distinctive mark of that thing’s institutional nature.)

Thus we have three basic ways in which the ontology of intentionality-dependent things can appear to us as more complex or obscure than the ontology of things whose basic ontology makes no reference to the intentional (“directed”) states of cognitive agents. We can refer to those three basic ways under the heads of *raw intentionality-dependence*; *dependence on collective intentionality*; and *dependence on status-functional forms of collective intentionality*.

I must point out that this classification is independent of discussions concerning the nature of intentionality itself (in other words, it is independent of whatever ontology of the mental we decide to endorse). Even if the belief that intentionality is the mark of the mental turns out to be false, and whatever “the mental” refers to (if at all), it seems to me that we can’t have an intelligible representation of the world without the idiom of agents’ “directed states”. We can’t talk intelligibly about the world and ourselves without the notion of “representations of the world”, beliefs which can be true or false, etc., whatever the correct ontology for things such as beliefs (and other directed states) may be. (It is a separate question whether talk about directed states should be paraphrased into talk about something else.) Should our “folk” ontology of the mental be proven utterly wrong, it would still be the case that the ontology of things naturally described as “intentionality-dependent” in the idiom of “folk metaphysics” is at least initially less clear. That is so partly in virtue of the fact that we tend most naturally to employ that idiom in describing the relevant things and events.

Ontological Contextualism and Ontological Historicism

Discussions of the relation between musical works and social reality also involve an extra layer of complexity introduced by “ontological contextualism” and “ontological historicism”, two approaches in the metaphysics of art which, as their names suggest, take, respectively, certain contextual and historical elements in the production of musical works (which may partially overlap) as ontologically determinative. These are determinative both of what general category musical works belong to, and of their conditions of identity (the properties in virtue of which any two sound events – assuming that being a sound event is a necessary condition for something to be an instance of music – are tokens of the *same work*). So, for instance, an ontological contextualist may hold that the general category to which musical works belong is in part determined by contextual properties of that work’s creation (say, certain facts about the musical tradition where it originally sprang). And an ontological historicist may hold the view that a musical work’s temporal career (the historical sequence of its embodiments, performances, interpretations, etc.) is an essential feature of its identity qua musical work. One example of overlap in contextual and historical properties is the application of the thesis of necessity of origin to the special case of musical works, since origin is both a contextual and an historical property of any given item.

Having said that, there is a considerable vagueness in the predicate “historical” which allows for a liberal interpretation, according to which all properties are “historical” (thus making the qualification “historical” void), since the exemplification of any given property is an *event*, and

events are the paradigm of temporally extended things – and there is no apparent reason to ascribe any stronger meaning to “historical” than “anything which takes place in time”. However, this liberal reading is not what the ontological historicist has in mind, in much the same way that not all historical properties of artworks are of metaphysical significance for someone who supports an historical (or recursive) definition of art. In this stricter sense, “historical” refers certain properties of any given item – call them *diachronic* – which contrast with that item’s *synchronic* properties. Before proceeding any further, let’s examine the notions of synchronic and diachronic properties.

I propose the following as working definitions of each kind of property.

- Diachronic properties: A property is a diachronic property iff it depends, for any given item that exemplifies it, on that item’s *actual temporal career*, over and above its being a (necessarily) temporal phenomenon.
- Synchronic properties: A property is a synchronic property iff the specification of that property makes no essential reference to actual “temporal careers”.

So, for instance, it is a synchronic property of the game of chess that it includes certain moves as legal moves of the game; it is a diachronic property of chess that it evolved out of Persian *chatrang*. There are possible worlds where the game of chess (or a functional duplicate of the game of chess) has a different temporal career; but there are no possible worlds where the game of chess (or its functional duplicates) have different rules (though there are possible worlds where the game, or a functional duplicate, has different recognized variants of the rules). One way of putting the contrast is the following: it is at least metaphysically possible that two people played chess (with the rules we all know) *before* any versions of the game (like *chatrang*) which actually lead to the introduction of modern chess were invented. The properties of the game for which this metaphysical possibility applies are synchronic.

Accordingly, *having been composed after Ravel’s composing of his **Concerto for the Left Hand** in D major* is an example of historical property in the narrow sense (diachronic), whereas *having been composed in the key of D major* is an example of synchronic property, even though such property depends on further historical properties – e.g. *having been composed after the development of western tonality* – the fact that the work is composed in the key of D major is not itself diachronic, which can be illustrated with the following example: a time traveler could go back in time and compose a work in the key of D major before the advent of western tonality, whereas he could not go back in time to 1905 and compose a work, in D major or otherwise, *after* Ravel’s composing of his work, which took place in 1929-1930. The time traveler’s composing of his work in D major before the advent of western tonality would still be causally dependent upon the latter

event and yet there is a sense in which the property of being in the key of D major is also a timeless or atemporal one. Another analogy: should our time traveler go back to a period before the coming into existence of any members of the *homo sapiens* species (of any tokens of the type), he would still be a member of that species, so that in a sense, his property of being human (a token of that type) is “timeless” even though it depends on a history of biological variation or “evolution”. By contrast, his property of having been born after the advent of western tonality is a diachronic property, since he cannot time travel himself into being born before he was born, the way he can time travel himself into composing works in the key of D major before the advent of western tonality.

With this in mind, the ontological historicist’s program is now made clearer: it is some subset of an item’s diachronic properties which have, for him, true ontological relevance. Some properties in that subset will also be contextual properties (diachronic properties of the work’s context of creation), and some will not. What distinguishes the contextualist from the historicist, then, is the emphasis (or absence of emphasis) in diachronic properties, since a context of creation includes both synchronic and diachronic properties. Take, for instance, the context of creation of the work we know as the *Seikilos Epitaph*: it is a synchronic property of that context that equal temperament was not conventionally adopted by the musicians working within the context (and indeed one could time travel himself into composing works for instruments tuned in the system of just intonation before the advent of just intonation); it is a diachronic property of that very context that it came out of the musical culture theorized three centuries earlier by Aristoxenus of Tarentum.

Now, take the metaphysical thesis of the necessity of origin (e.g. that *this* tree could not have originated from a different seed). This is a clear example of a historical property in the narrow sense (diachronic). In its application to musical works, the thesis of the necessity of origin gets somewhat more complicated, since “origin” here may include diverse elements, some of which will be synchronic properties of the work or its context of creation, and some of which will be diachronic properties of the work or its context of creation. One example of ontological historicism is the doctrine we may term “historical individualism” (Rohrbaugh: 2003), according to which musical works (and presumably artworks in general) are *historical individuals*, i.e., things which are essentially endowed with a “life” or a specific temporal career, with the distinctive trait that it is a feature of “historical individuals” so conceived that they are also temporally and modally flexible (they can suffer changes in their intrinsic properties over time and they could have had a different set of intrinsic properties from the one they actually have). The problem of the individuation of musical works (of what their conditions of identity are) is about determining what subset (if any) of

their intrinsic properties must be invariable across the possible worlds where the works exist (i.e., across the set of their actual and counterfactual histories).

The limitations of musical platonism

One of the salient features of musical Platonism is that it conceives of musical works as a-historical entities, i.e., entities which lack both temporal and modal flexibility – not changing over time and not having different intrinsic properties in other possible worlds in which they exist. Musical Platonism is distinguished from historicist ontologies not by the belief that musical works have no diachronic historical properties but by the belief that musical works do not suffer change across time – their temporal careers are inflexible – and that in no counterfactual temporal career could they be intrinsically different from what they are. And for the Platonist, only synchronic properties are relevant for a musical work's identity, which is another factor distinguishing him from ontological historicists in general.

What contextualists and historicists have in common is the emphasis on features of the social environment (synchronic or diachronic) in which musical works are composed, which means that both approaches take seriously a conception of musical works (and artworks more generally) as *social entities*, i.e., things whose ontology makes essential reference to some form of collective intentionality (in other words, their dependence upon joint action and shared representations is not merely a causal dependence, in the way an orchard may be the causal outcome of the joint action of individuals). Here the view of the institutional theorist differs from other approaches which underline social factors in that the forms of collective intentionality he takes as definitive for musical works (and artworks more generally) include, as we have seen, status-functional forms of collective intentionality.

If the theoretical approach, regarding both the definition of art and the definition of music I've adhered to in this work is correct, then these latter forms of collective intentionality are not essential for the existence of either artworks generally or musical works specifically, though, in a wide historical (but not historicist) perspective, they become more or less inevitable, since the same abilities that unavoidably make us producers of institution coincide (at least partially) with the abilities that make us producers of art generally and of music, more specifically.

One of my purposes has been to determine whether the problem of defining art as well as music sets any conceptual or metaphysical constraints on our answer to the problem of the ontology of musical works. My working intuition from the outset was that it does so. The exploration of the problem of definition served thus, in great measure, to test this intuition. It is now time to make

explicit my conclusions on whether the intuition turns out to be true and how exactly this can be shown to be true or at least plausible.

For a long time, the ontological debate has focused on the confrontation between musical Platonism, in its pure and qualified varieties (involving differences regarding the creatability of musical works, the role of the composer in the identity of musical works, and the issue of pure or modified acousmaticism) and musical nominalism – the thesis that apparent reference to abstracta in talk of musical works is in fact indirectly about concrete entities of one sort or another. A quieter companion in the “fundamentalist debate”, though present from the start (Rudner: 1950), is musical eliminativism – the thesis that names for musical works are non-referring terms, since there are no things which are musical works. Defenders of this theoretical stance are not always clear on whether the object of elimination is our *concept* of musical work or the referent of the concept (Pettersson: 2009). This lack of clarity should be avoided, since it may very well be that we *need* our concepts of musical works, even if it should turn out that those concepts really have no “targets” in the “external world”, no referents.

I am inclined against elimination regarding musical works for basically the same reasons I am inclined against elimination regarding artifacts generally. A ship is not merely a shipwise arrangement of “simples”, or wooden boards arbitrarily taken as “simples”. There is a level of conceptual dependence in the reality of all artifacts, whether causal-agentive or status-functional, intimately connected with the fact that artifacts are defined by (agentive) functional properties, which are not a mindless product of natural evolution. But conceptual dependence makes the dependent object as “unreal” as dependence upon an arrangement of “simples” would make it unreal. This could be seen as begging the question against the eliminativist – since the latter believes just this: that there are no ships, only shipwise arrangements of “simples”. However, what I claim is that conceptual dependence and dependence upon simples are ontologically on a par, and, in any case, the eliminativist does not believe that it is dependence upon arrangements of simples what makes the ship unreal. As I understand it, eliminativism about ships is motivated by the need to eliminate some puzzles about object coincidence. Since our conceptual activity is as much a part of reality as “simple” items are, I see no reason to tag any conceptually dependent items as “unreal”. Of course we can use the idiom of “real” vs. “unreal” to distinguish items which are independent of human conceptualization from those which are not; but that would beg the question whether “unreal” in *this* relative sense coincides with “unreal” in a non-qualified manner. No coinwise arrangement of metal particles is a sufficient condition for that arrangement to be *money*. In fact, the *same* coinwise arrangement of metal particles is money at one time and not (actual) money at a different time. I see no inferential route from this to “there really is no money”. Human

beings must exist within the boundaries of certain physical or biological constraints; a network of dependencies assures our presence in the world. I see no reason to deny reality to a domain of objects, should the network of dependencies that assures *their* presence in the world include the conceptual activity of human beings themselves.

In more recent times, the ontological debate has been enriched with the coming into play of contextual and historicist theories – an “abundance of riches” which is not an “embarrassment of riches”, since the paucity of conjectures about the nature of an X is no evidence that the conjecturing taking place has a higher degree of epistemic virtue than it would otherwise have; or at least I see no reason to think it so. The confrontation of a multitude of ontological perspectives can be dizzying, but it also provides a useful contrast allowing us a more fine-grained understanding of the problem’s contours, even if we are not able to arrive at a final, conclusive answer to the original question. I will not further press on this point, though.

As well as determining whether the definition of music sets any constraints on our theorizing about the nature of musical works (the problem of what ontological category they belong to), I’m interested in knowing whether the status of musical works as *socially constructed* entities has any bearing on determining their basic ontological nature and how. Evidently, this is connected with the question of whether the ontology of socially constructed things differs from the ontology of things which are not socially constructed. Hence my exploration of the forms of intentionality-dependence, for at least some forms of social construction are also forms in which the constructed items depend, in one way or another, on the intentionality of social agents.

Of course, there are many things which are the result of agents’ joint activity, activity which involves the sharing of beliefs as well as the representation of the several individual actions as parts of a larger collective action, but are nonetheless independent from the social nature of the action: an orchard or a building may be the result of several individuals working together, but there is no conceptual link between *being an orchard* or *being a building* and *being the result of collective action and shared beliefs*. The concepts “orchard” and “building” are intentionality-dependent in the relevant sense (namely, all concepts are products of our mental activity), but the objects picked out by those concepts, even though they depend on human action and as such on some form of mental activity, are not essentially the product of collective forms of intentional action. Such objects contrast strongly with entities such as a natural language, for language possesses all three layers of complexity we’ve specified, regarding the property of intentionality-dependence: raw intentionality-dependence, dependence on collective intentionality, and dependence on status-functional forms of collective intentionality. For instance, no utterance-wise arrangement of sounds or squiggles constitutes (counts as) an utterance in English merely in virtue of the acoustic/graphic properties of

the items arranged; it is not just about raw dependence on cognitive agents' "directed states", in the sense that a belief is only a belief if it consists in some kind of directed state – no *shared* directed states are required for a belief (a particular representation of some state of affairs in the world) to *be* a belief. Likewise, it is not just about dependence on collective intentionality, as in the case of actions which must be performed collectively, and there are two levels of representation, where the members of the group performing the action conceive of themselves simultaneously as performing some individual tasks *and* of themselves as part of a group acting as a single entity, i.e., the actions performed by the individuals in question are represented by each of them as "we do/perform X". Besides these forms of dependence on directed states of cognitive agents, language requires that the agents share some coordinated representations concerning what, and under which circumstances, counts as a correct or well-formed utterance of the language they speak. Whether a particular sequence of sounds/squiggles is to be taken as an utterance in any given language depends on certain contextual factors, which include the coordinated beliefs of speakers.

From this it seems clear that, in the case of languages, a pure Platonist ontology will not be able to capture the complexity introduced by such social and contextual factors. This observation is ambiguous between two readings, tied to the two-pronged question of musical work ontology: one as to whether pure Platonism is able to deal with socially constructed entities regarding the issue of the ontological category they belong to; another regarding the issue of musical work individuation. The two issues are distinct, though intimately connected, for what properties are relevant to musical work individuation is presumably constrained by the issue of what properties musical works, *qua* things of their *kind*, can and cannot have. Conversely, individuation may constrain the categorial problem: supposing the truth of the view that musical works *must* be individuated finely (e.g. including reference to composers and time of creation)²⁷⁴, then they *must* belong to a category whose instances are able to sustain such properties or have them co-instantiated in any given token. The direction of constraint here – whether it flows from the categorial issue to individuation or the other way around – is itself partly determined by what methodological constraints we choose to impose on our theorizing about musical works, or, to use the technical term of the art, our *meta-ontology* of musical works. Should we begin by determining what ontological category musical works belong to and only then proceed to make explicit their conditions of individuation? Or should we rather begin with individuation and only then proceed to extract the ontological category, having

²⁷⁴ The criterion of *fine individuation* was introduced by Levinson (1980) as one of three constraints on musical ontology, the other two being *creatability* (musical works must be creatable entities) and *instrumentality* (instrumental means of performance are integral to musical works' identity).

established the list of properties that are relevant for musical work individuation? Pure Platonists typically think that the categorial issue takes precedence over individuation, whereas contextualists typically consider that one cannot solve the categorial problem independently of what people (the grounders and re-grounders of art-kind terms) already believe about the way musical works are individuated, in which case the ontological category under which musical works are to be subsumed will be, so to speak, read off from the list of individuation-relevant properties. In this latter view, musical works *must* belong to whatever ontological category allows for its tokens to be individuated so-and-so (where “so-and-so” stands for the beliefs of a community of concept users).

In previous sections I have suggested that this picture generates a false methodological dichotomy between meta-ontological realism (the view that there are ontological discoveries to be made about art-kinds, and that our “nascent ontological conceptions” about such kinds are prone to error), associated mainly with pure Platonism and nominalism, and meta-ontological descriptivism (the view which denies the previous one), associated mainly with contextualist and historicist doctrines. This is supposedly due to the fact that the latter are highly prone to a revision of the “categorial map” bequeathed to us by general metaphysics: a map typically containing divisions between *universals* and *particulars*, *things* and *events*, et cetera. Contextualism and historicism have a tendency to break the boundaries fixed by traditional metaphysical schemes – e.g. considerations about fine-grained individuation may lead ontologists of music to postulate abstract entities which are creatable (thus contradicting the conventional wisdom according to which abstracta can’t come into existence nor cease to be, on account of their presumed non-spatiotemporal existence and thus causal inertness), temporally and modally flexible (thus contradicting the conventional wisdom according to which abstracta can’t suffer changes in their intrinsic properties over time, nor have their intrinsic properties contingently, i.e., differing in their intrinsic properties across possible worlds and not merely over time).

However, I see no reason why a revision of the conceptual or categorial maps or “conceptual schemes” typically adopted in the context of general metaphysics should not be motivated by developments in special metaphysics, in the case in point, the metaphysics of art kinds. And there is no compelling reason to suppose that doing so involves a compromise with meta-ontological anti-realism (usually associated with descriptivism), i.e., the methodological thesis according to which the nature of art kinds, and other socially constructed kinds, is *established* by the very act of their social construction. We may call this approach “social constructivism” about art kinds; its core idea being that the *content* of intentional acts of social construction *determine* the nature of socially constructed items, in particular the higher order *types* or *kinds* of which things like individual

artworks or individual works of music are the (second order) tokens, having in their turn sonic events – performances or otherwise – as *their* tokens.

Constructivism

I believe we do not have to endorse social constructivism about socially constructed entities (whether types or tokens) if we are to take seriously the idea that accounting for the nature of art kinds, and other socially constructed kinds, may require us to revise the typical categorial boundaries of general metaphysics. After all, it may be just a plain mistake to suppose that abstracta, if any such entities in fact exist, must be non-creatable in virtue of not having a spatiotemporal location; it may be a plain mistake to suppose that abstracta must be causally inert, etc. Considerations from special metaphysics (e.g. from art-kind ontology) are as respectful a source of reasons for such a revision as anything else.

Even if “folk ontology” or the “nascent ontological conceptions” in the minds of grounders and re-grounders of art-kind terms, concerning the nature of musical works, is to be our *prima facie* epistemic route to the nature of art kinds, this by no means makes it the case that social constructivism about art kinds is true. We may have nothing other than the *art practices* themselves to begin our enquiry into the nature of those entities but this is no evidence that practice, our *epistemic* route to a range of objects and events, should also be a sort of infallible guide in our metaphysical theories of such things. This would be so if meta-ontological anti-realism were true, at least concerning the special domain of art-kind ontology. This form of anti-realism would be true if social constructivism about socially constructed entities were true; yet no evidence in favor of this is to be drawn from the mere fact that artworld practice is our *prima facie* epistemic route to the items in question.

Social constructivism is not the only game in town for anyone with the working intuition that social kinds share ontological specificities distinguishing them from natural kinds. For any given item, dependence on complex social practices (Kania: 2008, pp. 76-77), on some form of collective intentionality (whether status-functional or not) is insufficient to establish social constructivism about that item as anything remotely like a default theory. An alternative approach would be to endorse a realist theory of social construction (such as the Searlean theory of social construction purports to be), in which case meta-ontological realism would hold for socially constructed kinds, as much as it holds for natural kinds, even if social kinds are to be distinguished from the latter in virtue of ontological specificities (e.g. intentionality-dependence, a feature not shared by natural kinds, even though it is evidently shared by natural kind *terms* or *concepts*).

To say an item *ontologically* depends on another is not the same as asserting a relation of causal dependence of the former towards the latter. We may call this second form of dependence *existential dependence* – whereby an item depends *for its existence* (either its coming into existence or its continued existence) on some further item which *causes* the former to exist. Existential dependence can also assume non-causal forms; for instance, sets depend existentially, but not causally, on their members – the set exists *because* its members exist, but its members do not cause the existence of the set, in the way the actions of a gardener cause the existence of the garden, or in the way the existence of an organism is dependent on the existence of previous organisms that generate it (i.e. that cause its existence). Ontological dependence, in the sense I’m deploying that term, is neither causal existential dependence *nor* non-causal existential dependence, but something stronger: an item X is ontologically dependent on an item Y if the *category* (in our “conceptual map” of reality) to which X belongs is determined by some feature of Y (e.g. something about Y *makes it the case* that X is a concrete particular object, rather than something else). The special brand of ontological dependence that pertains to social constructivism about socially constructed items takes on the following form: the category (in our “conceptual map” of reality) to which the socially constructed item X belongs is determined by the beliefs of social agents Y, who “socially construct” X, rather than X belonging, in a mind-independent manner, to a certain ontological category, to which we *refer* using certain concepts and words.

There is another kind of dependence relation, stronger than causality but not as strong as ontological dependence (in the sense I’m deploying that term): that relation is sometimes referred to as a relation of *supervenience* or *co-variance* between properties of an item. A property or cluster of properties B supervenes on a property or cluster of properties A iff two items, X and Y, cannot differ with respect to their B properties unless they differ with respect to their A properties. We have already seen how this concept of supervenience or co-variance applies to aesthetic properties: when we say that aesthetic properties depend (supervene on, co-vary with) on certain non-aesthetic properties, that means not that the presence of certain non-aesthetic properties in an item X causes the presence of certain aesthetic properties in that item (in the way that the presence of an infection in an organism may cause fever) but something quite different: we say that a certain contrast of colors *makes* the harmony of a painting, but not that the contrast *causes* the painting to be harmonious; since *harmony* is not just another of the painting’s perceivable properties; strictly speaking, what we do *perceive* is the contrast of colors, the harmony (or lack of it thereof) being a quality the painting has *by way of* having that contrast of colors. If we strip an object of its perceivable properties, there are no aesthetic properties left for us to be aware of. When two

properties are causally related, (at least typically) one can be present in an item without the other being present; co-presence and co-variance are not necessary conditions for the causal relation to obtain.

Supervenience is thus a dependence relation somewhere between causation and identity: the harmony of the painting is not caused by its contrast of colors, but neither is the harmony *identical with* that contrast. This can easily be seen by the fact that the same chromatic contrast may occur in another painting without it being the case that the painting is harmonious – other factors may intervene to prevent our experience of the painting to be the experience of a harmonic whole. The definition of supervenience merely states that any two items can only differ with respect to their B properties iff they differ with respect to their A properties (e.g. that two items, X and Y, can only differ aesthetically if they have some non-aesthetic difference). The class of A properties on which an item's B properties supervene is what we call that item's *base of supervenience* or *subvenient properties*. Even among philosophers who endorse the thesis of aesthetic supervenience (the thesis that aesthetic properties supervene on non-aesthetic properties) there is disagreement about how wide or inclusive the base of supervenience is. But such disagreement does not alter the nature of the relation being postulated between the two classes of properties.

Having said that, the relation of ontological dependence (in the sense I'm deploying that term) is not a relation of supervenience either. The property a socially constructed item X has of belonging to a certain ontological category does not supervene on the propositional attitudes of the social agents Y, who socially construct the item X. That is not the proper kind of relation obtaining between these two sets of features. If people believe musical works to be creatable items, and if the social constructivist is right, then musical works, if they exist at all, must be things which have the property of being creatable; but the property of being creatable is not something which co-varies with people's propositional attitudes about creatability, since many items are creatable without it being the case that they are such in virtue of anyone having certain propositional attitudes.

The kind of relation the social constructivist believes to obtain between the ontological nature of socially constructed items and social agents' propositional attitudes about those items must be a relation of determination which is neither existential (causal or non-causal) nor of supervenience or co-variation, but of a different kind. According to the social constructivist (in the sense I'm employing this term), the social agents who "construct" the items in question determine their ontological nature (i.e. the ontological category to which they belong in our "conceptual map" of reality) by *making it the case* that socially constructed item X belongs to category C, in virtue of their belief that X is of kind K (a kind that falls under C). This relation of *making it the case that...* is not in the causal sense of *to make* [an X], where the existence of an X is the causal outcome of

the agent's productive efforts (which include insights concerning relations of properties and planning). This is a more queer case of determination of the metaphysical by the epistemic: (presumably) we make it so by thinking it so. It is "queer" because the trivial cases where *thinking* is sufficient for *making* are not cases where the ontological category of the dependent item is determined by the content of our directed states toward that item, in this case, beliefs.

Now a question suggests itself: besides socially constructed entities, are there any items of a kind K, falling under category C, where the fact that an item X is an instance of K, falling under C is itself a consequence of some cognitive agent Y (here, whether the directed states are "social" in character is not relevant) representing X as of kind K, falling under C? I can think of no examples. Even in the case of divine creation, or creation of an X by an omnipotent, all-knowing being, it seems clear that it is not the *representations* of the creator in creating, say, a biological organism, what makes it the case that the created item falls under the category *concrete particular*, or that it is a four-dimensional entity (temporally as well as spatially extended), etc. A divine creator could choose *what* to create, but cannot make it the case that reality is not exhausted by the division between particulars and universals (if there are any). That a divine creator would be thus limited is a stance with a long history in the philosophy of religion (e.g. in Aquinas's response to the so-called "paradox of the rock"), where God, despite his omnipotence, is conceived as not being able to generate logically contradictory states of affairs (such as creating a rock so heavy that even he could not lift). What this means is that even a divine creator would be constrained by logical necessity (a feature of objective reality) – e.g. the divine creator would be unable to change the ontological category to which *he himself* belongs, neither could he – should the world be such that it contains a divine creator – make it the case that the world is not such that the existence of a divine creator is one of its possibilities.

But let us leave the analogy at that point. If socially constructed entities are the only cases the social constructivist can point out, of entities for which that particular relation of determination obtains, it seems to me that his case becomes somewhat more fragile, though this is not, taken by itself, evidence that social constructivism (in the sense I'm deploying that term) is false. It merely lacks the extra strength or plausibility it would draw from the analogy with other kinds of entities. But one could not object to the view exclusively on that ground, for it would beg the question against the claim that socially constructed entities are ontologically *sui generis* in that regard (this *might* very well be true).

Earlier I claimed that how far we move from a pure Platonist ontology of musical works towards contextualism or historicism is correlated with the relative weight of our concerns for individuation and the categorial question (an emphasis on individuation over the categorial problem will most

likely yield an ontology of art more open to a revision of traditional metaphysical schemes and more sensitive to contextualist and historicist qualms). This is immediately connected with a methodological issue: how should the musical ontologist proceed? Do we *begin* by determining the ontological category that musical works must belong to (categorical problem) and *only then* move to the question of what makes any given work the work it is (individuation)? Is it *possible* to solve the categorical problem before we try to solve the individuation problem?

I suggested there is not a clear dividing line here between the two approaches, but that they rather stand in a continuum, depending on the list of properties one considers relevant for the individuation of musical works. If we put things this way, it becomes apparent that even a pure Platonist such as Dodd does in fact begin with individuation. How is that, if in his book on the ontology of musical works, Dodd (2007) devotes the first half of it to the categorical problem, only then moving on to the problem of individuation? The key to understand this lies in the way he states the categorical problem:

“What kind of entity must a work of music be, given that it can have multiple occurrences? And how, given that an occurrence of a work is distinct from the work itself, is it possible to listen to the work in listening to an occurrence of it? An ontological proposal for works of music should try to answer these questions.”²⁷⁵

The problem is put in terms of explaining certain metaphysically perplexing properties of musical works, namely, musical works are *audible* and *repeatable* (they can have multiple instances). So the categorical problem is stated in terms of determining what kind of entity can be both hearable and repeatable, in the sense that certain spatially discontinuous (whether simultaneous or not) *events* may count as *occurrences* of one and the same entity; more specifically, since several spatially discontinuous *sound events* (simultaneous or otherwise) are in fact counted as occurrences of one and the same musical work, what kind of entity must a musical work be if we are to take seriously the claim that musical works are both hearable and repeatable?

When it comes to the issue of individuation, Dodd defends *sonicism* against its immediate rivals: instrumentalism (the thesis that instrumental means of sound production are part of the individuation criteria for musical works) and contextualism. Sonicism is the doctrine according to which the only kind of properties relevant for the individuation of musical works are, as the name indicates, *sonic* or acoustic properties. This brings us back to the list of putative properties of

²⁷⁵ Dodd (2012: p. 1052).

musical works that Dodd (and any musical Platonist) considers “metaphysically perplexing” and thus requiring explanation (the explanatory purpose of ontological theories being the major reason Dodd rejects descriptivism in the metaphysics of music): hearability and repeatability. We can further state *what* is so perplexing about these properties in terms of what it takes to co-instantiate them in a single entity. The problem lies in the *compatibility* of these properties qua simultaneous properties of one and the same thing; whereas *hearability* seems to be the kind of property that pertains to *concrete entities*, namely, *sounds* (which are arguably a kind of *event* and thus belong to the class of temporally extended concrete items), *repeatability* seems to be the kind of property that pertains to *non-concrete entities*, since no paradigmatic concrete entity (objects and events) are repeatable, i.e., none is the kind of thing that can have *instances* or exist wholly in discontinuous regions of space and time. Two sound events may belong to a single kind, e.g. they can both be the *sounding* of bells, but neither event is in itself repeatable; once a sound extinguishes, it is no more; the sounding of a relevantly similar sound (e.g. a timbrally identical sound) is not the same sound. Two sounds, no matter how qualitatively similar, are always *numerically distinct* events. This prompts the question: what kind of thing is both hearable and repeatable? Such is the framing of the categorial problem for the musical Platonist.

Looking close at the properties in terms of which the Platonist frames the categorial problem, we see the connection with the problem of individuation: once settled on the most plausible ontological category available under which musical works, described as hearable and repeatable, may fall (e.g. non-concrete or abstract *types* or universals), and after telling a story aimed at explaining how these properties can be mutually compatible, our musical Platonist will ask whether *any other properties*, besides sonic properties, are relevant to determine in what circumstances a particular sound event can be said to instantiate (to *repeat*) or to be the occurrence of one and the same musical work. In other words, the properties which are relevant for musical work individuation and the properties in virtue of which musical works are metaphysically perplexing are basically the same. To be sure, even where artworks are *prima facie* identified with singular concrete objects, as in the case of paintings, the issue of individuation arises (it is not entirely clear what makes any given singular artwork be the artwork it is across times and possible worlds), so that repeatability is not a special source of perplexity for individuation as it is for the categorial problem: with any multiple (repeatable) artwork one must ask what properties any token must instantiate in order to count as a token of *that* work; with any singular (non-repeatable) artwork one must ask what properties any temporal stage of an object, bearing the appropriate causal connections to an earlier temporal stage which we identify with the original work, must instantiate in order to count as a bona fide stage in the temporal career of the original work. In both cases we are concerned with a minimum cluster of

properties that something (whether an *instance/occurrence* or a *temporal stage*) must instantiate; that cluster of properties will comprise what we call the *criteria of individuation* for the kind of work in question.

To the extent that repeatable entities are most readily thought of as non-concrete (or abstract), and to the extent that such things are commonly thought of as not having spatial location and thus causally inert (causation being commonly thought of as a relation between things which have spatial parts capable of physical interaction), repeatability seems to bear on the individuation of musical works across possible worlds: if the most common view of non-concrete entities among metaphysicians is right (and if such things do exist) then musical works (qua repeatable entities) could not have different intrinsic properties from the ones they actually have. Yet, this problem also arises for singular works which are thought of as concrete physical entities: it is not entirely clear whether the object we identify with the work could have *certain* different intrinsic properties from the ones it actually has (would the *Mona Lisa* be the numerically the same painting if it had been painted with a different batch of qualitatively indistinguishable paints?), and so it is not an issue which becomes especially complex in virtue of repeatability. The most salient difference between singular and multiple works in this respect is the fact that the temporal stages of singular artworks are (at least typically) continuous and causally interconnected, whereas the instances or occurrences of musical works (and repeatable works generally) are (at least typically) spatiotemporally discontinuous and not necessarily causally connected (two causally disconnected performances of the same musical work may both be causally connected to the original score by way of a chain of typographical reproductions and readings). Given that difference, we can think of singular works' temporal stages by analogy with repeatable works' instances/occurrences, with the above mentioned difference concerning spatial and causal continuity between the temporal stages of singular works. Apart from that, individuation of musical works must deal with property variation across performances (instances/occurrences), analogously to the way individuation of singular works must deal with property variation across temporal stages. Of course, even if singular artworks were rightly conceived as inflexible in certain of their modal properties, it would still be true that they are temporally flexible: they can suffer change in their intrinsic properties over time. For instance, as a painting ages, its paint cracks and exhibits changes in color, without it being the case that the painting now (this temporal stage of an object) is a numerically distinct painting from the painting (a past temporal stage of the same object) to which it bears the appropriate causal connections. But if a painting can suffer such changes without ceasing to be the *same* painting, then it could suffer different changes over time from the ones it actually did without thereby ceasing to be the same painting. Thus, if an artwork (or any other item) is temporally flexible, then it must also be modally

flexible (unless the unappealingly counterintuitive theory that the actual world is the only possible world was true). Nonetheless, whether or not a given artwork is temporally flexible (and thus modally flexible), if it has conditions of individuation at all, then there is a subset of the artwork's intrinsic properties which cannot vary without making it the case that the artwork is no longer numerically identical with whatever results from that variation in properties.

Bottom-line, the musical Platonist, no less than the ontological contextualist and the advocate of "historical individualism" about musical works, theorizes about the categorial problem equipped with a list of properties that all things which are musical works must have (according to our musical Platonist, nothing is a musical work which is not hearable and repeatable). These are the same properties the musical Platonist considers relevant for musical work individuation, namely, sonic or audible properties. In other words, no one begins to consider the categorial problem with musical works completely devoid of properties (if such an exercise in abstraction is even possible). No bare particular (if such a thing is even conceivable) can be a musical work; to be a musical work *is* to have certain properties rather than others (to begin with, it is to have the complex property of *being a musical work*, which will itself consist in the possession of further properties); the categorial problem, as much as the problem of individuation, consists in determining *what* properties these are. Though both problems are distinct (it is one thing to ask what makes any given item a *musical work*, and another to ask what makes any given musical work *the musical work it is*), they are always connected by the fact that whatever properties are responsible for a thing's being a musical work are also properties that thing cannot lose on pain of not being the musical work it is. We can put this point in terms of the difference between *essential properties* and *essential individutive properties*. Not all essential properties of an item are individutive properties of that item, though all individutive properties of an item are also essential properties of it.

Given that all individutive properties are also essential, any property which is thought of as individuation-relevant will be categorially relevant as well, since no item can lose an essential property without ceasing to be what it is. Take the case of biological species. If it is true that no swamp-duck (an organism which randomly and miraculously came into existence when a lightning stroke a swamp and which is a functional duplicate of a member of the Anatidae family) is actually a duck – because, in order to be a duck, an individual organism must stand to other individual organisms (members of the Anatidae family) in certain *historical* (in this case, evolutionary) relations, namely, it must share with them a *line of descent*, in order to be a bona fide duck – then such historical relational properties are categorially relevant; it is not enough to place ducks (and other biological organisms) within the category of concrete particulars, for this falls short of distinguishing biological organisms from other entities in the realm of concrete particulars (namely,

swamp-ducks). Furthermore, even for those who feel the need to postulate an abstract type or universal to explain what binds together all ducks (namely, that they are tokens of the type *duck*), the historical relational properties (lines of descent) will figure among the properties that any token biological organism must instantiate in order to be a bona fide token of the relevant type.

In fact, any addition of properties to the list of putative individuation-relevant properties for musical works will likely require a tweaking in our approach to the categorial problem; it may lead us to postulate further ontological categories than the ones we inherit from the typical categorial maps of general metaphysics. But such revisions to our conceptual scheme need not be radical (individual biological organisms conceived as essentially historical entities are still concrete particulars, and putative “historical types”, whose tokens are “historical” concrete particulars, are still types) nor need they be ad hoc additions to the scheme – after all, a “categorial map” which includes *events* alongside concrete objects which are not “extended in time” must be the result of taking into consideration that crucial difference in the individuating properties of a subclass of concrete particulars (events). The general metaphysician theorizes *on what there is*, and special metaphysics (which must inevitably be empirically informed) is as much a source of information on what there is as any other domain of enquiry.

It would seem that a methodological guideline recommending that emphasis should be put on individuation over the categorial problem favors metaphysical descriptivism over revisionism (i.e. revisionism of traditional “categorial maps” from general metaphysics over faithfulness to these); but that is in fact a conceptual illusion, perhaps rooted in a confusion between metaphysics and epistemology. Whence such confusion? Given that metaphysical descriptivists about musical works claim that, for socially constructed entities, epistemology *determines* ontology (e.g. Thomasson), this may not be fairly described as a *confusion* but a plain philosophical mistake (descriptivists are not confused but simply wrong). The root of the mistake (or confusion, if there should be any) is to think that *because* individuation-relevant properties help us track categorially relevant properties of musical works (and artworks generally), *and* since we have no other source for what properties *might* be individuation-relevant other than art practices themselves (which crucially include critical and appreciative practices), then the representations (e.g. beliefs) which constitute common wisdom within those practices cannot include (gross) *misrepresentations* (or errors) about the nature of musical works (or artworks generally). We have already seen that this idea can be challenged, and that even if it were true, it would be compatible with meta-ontological realism; in fact, it would *have* to be compatible with such realism, since the thesis that shared epistemic states of social agents *determine* the ontological nature of socially constructed entities is the true cornerstone of social constructivism about socially constructed entities (descriptivism being merely a

methodological concomitant of such *metaphysical* doctrine of determination). If the determination thesis was true, then the relations of determination that it represents as obtaining between social agents and socially constructed entities would be a matter of objective ontological fact, independently from the epistemological issue of whether or not the shared beliefs of social agents are *truth-tropic* (i.e. whether they *track the truth*) concerning the objective ontological facts about socially constructed entities. In other words, the property of being ontologically determinative and the property of being truth-tropic are both *conceptually distinct* and *metaphysically independent* properties. Furthermore, the property of being truth-tropic *comes in degrees*, which is not the case with the property of being ontologically determinative (such a relation of determination either obtains or it does not). As such, the presence of a degree of the former is compatible with the falsity of the thesis of epistemic determination of ontology for socially constructed entities. In fact, even when theorizing about natural kinds (for which the determination thesis intuitively does not apply), one must start with the way cognitive agents divide reality into categories; it is simply not true that we arrive at a metaphysical (or, for that matter, an empirical) description of any given natural kind without it being the case that the previous epistemic states of cognitive agents (our “folk metaphysics”) function as a guide in our enquiry on what there is. To that extent, those epistemic states (including beliefs and other representations) are also, to some degree, truth-tropic concerning how the world divides into natural kinds. The presence of that property does not imply any relation of determination of the kind the social constructivist postulates between social agents and socially constructed entities. And even upholders of metaphysical descriptivism recognize that there are degrees of inconsistency in the representations to be found in our art practices, inconsistencies to be trimmed out by “rational reflection”(see D. Davies’ formulation of the Pragmatic Constraint), so that even here the property of being truth-tropic is only present to some degree in the beliefs (and other representations) of the relevant class of “social agents” (e.g. artists and art critics).

We are now in a position to state the intermediate link between art-kind ontology and general metaphysics, enabling us to “break the circle” between one-sided “metaphysical fundamentalism” and an equally one-sided descriptivism, hanging on a dubious ontological determination thesis. That intermediate link is the metaphysics of socially constructed entities, of which things falling under art kinds are a subclass. This is the place to enquire whether the determination thesis (the metaphysical thesis on which social constructivism must draw if its scope is to be more than merely epistemic) is true or whether it must be true for all (or any) kinds of socially constructed item. Our adventures into the realm of social kinds and institutional kinds strongly suggest (at least) that nothing about social construction entails social constructivism. Certainly metaphysical descriptivism, which is an *epistemic* thesis, does nothing to establish the determination thesis on

which social constructivism depends. Descriptivist qualms are compatible with meta-ontological realism: we must in fact begin our theorization with *some* list of properties that *prima facie* characterize the entities which interest us; no such list is available to one who merely considers the traditional “categorical maps” of general metaphysics, without some delving into the representations of art kinds within artistic practice. But this does not mean that theorization begins and ends with an empirical scrutiny of our practices. Furthermore, given the wide range of things that fall under “socially constructed item”, there is no Pragmatic Constraint to help us sieve out rival theories. There is no “folk social ontology” to appeal to, and even if there was, it would constitute data to be considered, not a limit or constraint on metaphysical theorization.

We must “look and see” how musical works are individuated (in the epistemic sense of “individuated”) within our artistic practice if we are to approach, however roughly, some insight about the categorial problem for musical works. But to look and see is not enough; it cannot give us the whole story. Epistemic individuation may be our first guide to metaphysical individuation, but it is no surrogate for the latter.

It is now time to end this section with a consideration of how social practices figure in the ontology of musical works; how an appeal to social practices may actually throw light on their ontological nature.

The consideration of the role social practices play in the ontology of musical works brings with it a consideration of the extent to which contextual and historical factors are both individuation-relevant and categorially relevant.

Initiated types and indicated structures

An important contribution to the clarification of these issues has been given by Robert Howell in his paper “Types, Indicated and Initiated” (2002). There, Howell sets up some key distinctions between *properties*, *patterns* and *types*. In his view, combinations of properties determine the existence of patterns, which in their turn (partially) determine the existence of types. But the existence of types (at least those types which constitute socially constructed items), according to Howell, is also mediated by the *actual uses* that patterns are put to within the context of a social practice. Thus, in order to *count* as a constituent of a type, a pattern must acquire certain *functional properties*; it is not enough that the pattern itself exists, in virtue of the combinations of properties that constitute the pattern being genuine possibilities concerning the way the world is (i.e., it is not enough that the world be such that it contains those patterns as genuine possibilities). Thought of in this way, types appear as *time and context tethered entities*, and thus as *temporally initiated* items,

things that come into existence and may cease to exist, unlike the patterns that partly determine them, and the combinations of properties that determine the patterns. Though not all types thus conceived arise out of human practices, many in fact do. But even (at least some) natural types, in this view, share with humanly made types the property of being temporally initiated. One such example is that of biological species. Howell writes:

“Properties presumably now exist that detail various genetic and protein combinations that specify many biological patterns besides those that belong to actual species such as the badger, the wolverine and *Escherichia coli*. (These include patterns that cannot now be instantiated, given current physical conditions in the universe.) Yet no one claims that, besides the current actual species, all these patterns constitute a myriad of other species that now actually exist – that now actually are present in the real world – but happen to be uninstantiated.”²⁷⁶

“Natural types appear to exist only when the patterns that they involve play a role in actual causal chains in nature.”²⁷⁷

This claim makes a strong stance against Platonism about types, according to which all types are eternally existing entities, both temporally and modally inflexible. Howell locates the reason why the Platonist thinks of types in this way in the latter’s conception of how the existence of properties relates to the existence of types. His remarks and criticisms are directed specifically at Julian Dodd’s earlier writings on the subject of musical ontology (2000; 2002), but they apply equally to any “strong” Platonist view (one which conceives musical works not only as abstract types, but as eternally existent entities as well). According to Howell, the belief that types are eternal existents is tied to three claims: a) that all properties exist eternally, b) that all types *K* are identified by a “property-associate”, that of *being a* (properly formed) *K*, and c) that all properties are property-associates of some type.

Howell rejects claims (a) and (c) while maintaining claim (b). Against the view that all properties exist eternally, he argues that certain properties involve “essential relations to contingent, temporally initiated entities” (such as the property of being an Elizabethan playwright, or of being a son of Abraham Lincoln); and as such they cannot exist eternally but rather come into existence only when the contingent entity or entities that play the role of *relata* in that relational property come into existence themselves. (He has no need to defend the further thesis that *no* properties exist eternally; for his purposes, it suffices that *some* properties do not exist eternally). If Howell is right

²⁷⁶ Howell (2002: pp. 116-117).

²⁷⁷ *Ibid.* p. 123.

and it is not the case that all properties are eternal existents, then neither can the types of which those properties are property-associates be eternal existents. A type which is identified by a non-eternal property-associate cannot itself be an eternal existent. (E.g. if the property *being an Elizabethan playwright*, which identifies the type *Elizabethan playwright*, is not an eternal existent – since it involves essential relations to contingent entities, namely, a socio-historical context – neither will the type be an eternal existent.)

But the case against Platonism does not rest solely on the idea of there being at least some non-eternal or *impure* properties. As I've mentioned, Howell introduces the notion of a *pattern* (already deployed by Kendall Walton (1988)²⁷⁸ – “The Presentation and Portrayal of Sound Patterns”), patterns being determined by combinations of properties. In this way, patterns function somewhat as intermediaries between properties and types. In the case of socially constructed types, mediation is completed by acts of *selection* and *indication* of the relevant patterns (and acceptable variations thereof), which accounts for the intentionality-dependence of such kinds. In the case of at least some natural kinds, mediation is completed by “causal chains” responsible for the selection mechanisms which give rise to “natural functions” (e.g. in the case of biological organisms). Mediation must be completed either by the relevant directed states of human (or other) agents, or by natural selection. But intentionality-dependence here is in itself no more mysterious than dependence on natural selection mechanisms and does not involve any determination of ontology by the epistemic states of social agents. No act of selection or indication by an agent can change the objective truth about the ontological nature of socially constructed kinds, though it *can* change *what* portion of reality one is referring to (given that different concepts may be expressed by more than one token of the same word-type, such as “work”). That the concept of “work” operative in a certain period of music history is not operative for *all* times within that history or across musical traditions does not lend any strength to the descriptivist's determination thesis. Musical ontology may be a more diverse landscape for other reasons: not only can there be musical entities other than “works”, but those things properly described as “works” may be diverse and yet no less objective than mountains, even though what identity criteria we *choose* to incorporate in our concept “mountain” may have an impact on the range of entities we refer to when we utter the word “mountain”.

²⁷⁸ Already in that paper Walton suggests that a pattern is more than the mere combination of properties, since we can distinguish different patterns in one and the same combination of properties. It takes human intentionality to isolate the relevant pattern out of the myriad possibilities. He provides the simpler example of a chequered pattern in a cloth and the diverse “readings” of it.

One important contribution of Howell's paper was to clarify the role that a socio-historical context plays in the individuation of more than one kind of socially constructed entity, his main examples being those of words and other symbols (such as the hooked cross in both pre-Columbian Amerindian cultures and in Nazi Germany), and of musical works. Take the case of the hooked cross: it exemplifies a geometrical pattern (a combination of properties) which figures as an intrinsic property of two distinct *types*: a symbol of fire in pre-Columbian cultures, and, from 1920 onwards, a symbol of the German Nazi Party. Now, if each symbol is identical with a numerically distinct type, and they both share the same geometrical pattern, then the geometrical pattern cannot be what *individuates* each symbol-type, even though it figures as an essential property of each, a property no concrete item could lack under pain of not being a token of any of the two types. Applying here the concept of property-associates of types (which Howell takes from Dodd), we would say that the combination of properties specifying the relevant geometrical pattern is *insufficient* to constitute either of those two more complex properties which are the property-associates of each separate symbol-type (the pre-Columbian symbol and the Nazi symbol), each expressible in formal terms as *being a K* (where *K* stands for the individual type in question). This means that the geometrical pattern must acquire yet another essential property (a property no concrete item can lack under pain of not being a token of the type in question) in order to constitute a *symbol-type*, rather than a mere geometrical type (say, the type *hooked-cross-wise-figure*), determined solely by the combination of properties that specify the pattern. From this we gather that symbol-types, unlike pure geometrical types, require *for their existence* the relevant directed states of social agents. This requirement is not that of mere causal dependence, in the sense that social agents' thoughts and actions are required to "kick-start" symbol-types into existence; rather, the property of *actually being used by a social group* (Howell uses the term "community") to carry the relevant additional properties (in this case, *semantic* properties) is an *essential property* of the symbol-type qua *symbol-type*, a property no item can lack under pain of not being a token of the relevant type. It takes the geometrical pattern actually being used by a social group in a certain way (and, we should add, in a particular context) to make it the case that the specific symbol-type *exists*. Using a well-known language in order to make this point clearer, we can say that the type *supervenes* on these three "base properties": (a) exemplifying the appropriate pattern *P*; (b) actually being used by social group *S* in a certain way; and (c) being so used within the appropriate context *C*. Putting things this way even helps us to make clear in what way (at least some) socially constructed things are said to be socially *constructed* (and essentially so).

This is a more detailed exploration than the one Howell provides in his paper, but it does not take a great leap of imagination to see that the structure of the dependence relation between the symbol-

type and the three “base properties” closely follows Searle’s formulation of status-function declarations: *X counts as Y in C*, where *X* stands for any object, event or agent, *Y* stands for the relevant status-function, and *C* stands for the context within which the *operation* of counting tokens of *X* as tokens of *Y* is enabled or becomes effective. To the extent that any item for which *X* is a placeholder will exemplify a potentially relevant pattern or combination of properties, both Searle’s and Howell’s “conceptual schemes” may be combined to generate a more complete and clear specification of the ontology of socially constructed kinds. The importance of this is directly proportional to the fact that Searle’s theory of social construction is presented as an *ontology* of institutional reality, even though an answer to the categorial problem concerning socially constructed kinds is nowhere to be found in Searle’s theory. In fact, he seems to deliberately ignore any consideration of that issue, as if the categorial problem was the consequence of a sort of conceptual illusion: as Searle describes them, “monism”, “dualism” and “trialism” – the views according to which, respectively, there are one, two or three *kinds of being* or basic domains of reality (namely, the physical world, the mental world and the Frege-Popper “third world” of abstract entities) – are all misguided in their attempt at counting how many “layers” make up the metaphorical layered cake we call “reality”. Monists start counting and stop at one, dualists stop at two, and so forth. According to Searle, the mistake is to start counting altogether. For all practical purposes, he is a monist, since he believes the several “layers” are but “levels of description” of what is basically *one and the same* reality, one and the same world. Yet this description is ambiguous enough to allow for several interpretations: it is still the *same world* whether or not it contains abstract entities (universals or types). However, a view of types where either *talk* of types is conceived as an oblique *façon de parler* about concrete particulars, or a view of types as themselves concrete (e.g. Eddy Zemach’s sui generis nominalist theory of types), are both compatible with the view that we are always dealing with different “levels of description” within one and the same basic reality. This means that Searle’s framework is compatible with different ontological approaches and seemingly none of his qualms about counting the cake’s layers is sufficient to convince us to abandon the categorial problem as fruit of a conceptual illusion.

On the other hand, if the Searlean approach to social construction and institutional reality is correct it *does* provide us with a list of essential properties which are conjunctively sufficient to individuate socially constructed kinds: for any given social kind, what makes any concrete item a token of *that* social kind and not another is the conjunction of the three properties already mentioned – (1) that the item instantiates a certain pattern or combination of properties (2) selected by a rule or a collective decision as required for the performance of a certain function, (3) the imposition of which requires in its turn an appropriate context. But in providing us with such

conditions of individuation for any given social kind, when we consider each specific pattern, function and appropriate contexts, the general form of a status-function declaration (*X counts as Y in C*) also provides us with a sketch of what kind of properties will be essential properties of *any* social kind which is also an institutional kind: there is always an operation of social agents' directed states onto some pattern (whether that pattern belongs to an object, event or agent), which always results in the imposition of a function, and all this is unexceptionally enabled by an appropriate context, whereby both the item instantiating the required pattern and the function imposed connect to other parts of the social world, namely, other institutional entities (which *constitute* the relevant context). The point about context is generally expressed in Searle's theory by the ideas of an indefinite (vertical) *iteration* of the status-function declaration (an item *X* on which a status-function *Y* is imposed becomes the *X* term of another application of the same formula) and of an indefinite (horizontal) *interlocking* of social and institutional facts. For instance, a certain pattern or a certain variation-range of the same pattern *counts as* composer *C*'s work, but this fact interlocks with the further social fact that, say, equal temperament is the tuning system standardly adopted within the community or social group in which composer *C* *indicates* the sound pattern which becomes an essential property of his work. Of course composers can compose works for different tuning systems than the one which is standardly adopted – we have already mentioned Harry Partch's and Lou Harrison's use of just intonation – but it is precisely that *contrast* which endows such works with many of their artistically relevant properties. Arguably, Lou Harrison's *Serenade for Guitar and Percussion* (1978) would not be the *same* work, were it written in and for a context where equal temperament was not the standardly adopted system. This is because the *contrast* with equal temperament is part of what gives that work its peculiarity or its character. A view of the ontological nature of musical works which takes only an abstract pattern of sounds as both categorially and individually relevant has no way of accounting for the role such contrasts play in determining the character of works.

We thus see, combining the approaches of Searle and Howell, how a connection with the directed states of human (or other) agents and, via these, with a historico-social context, may figure essentially in a description of the ontological nature of musical works (and other socially constructed entities). But now a doubt arises concerning those artifacts which are not status-functional – i.e. causal-agentive artifacts, those which do not owe their status to an application of the rule *X counts as Y in C*. How do *these* artifacts fit the scheme generated by the conjunction of Searle's and Howell's views? This question is important, since one of our previous conclusions was that the status of an item as either an artwork or a musical work is ambiguous between a causal-agentive concept and a status-functional concept, in such a way that the latter depends or is

subordinate on the former, namely, it is because we have a causal-agentive artistic or musical function that we are able to create and impose the status-functional correlates of these.

The relevant directed states of human agents that “operate” on patterns (combinations of properties), as we have by this time abundantly seen, are of basically two kinds: (a) insights about dependence relations between “base” physical properties and functional properties sustained by the base properties; (b) shared representations about functional properties that a given item may perform *not solely (or not at all) in virtue of its physical structure but in virtue of those shared representations*. As we have also seen, these two kinds of relevant directed state mark the difference between causal-agentive artifacts and status-functional (institutional) artifacts.

Now, the question is: does it make a difference *ontologically* whether the artifact is causal-agentive or status-functional? Does it make a difference for the way Howell’s useful distinction between properties, patterns and types apply to *all* artifacts (if it does) and not merely to status-functional ones?

Partly, the answer to this question is tied to another one, having to do with the *ontological diversity* within the class of musical works, or rather with the question of whether or not such diversity exists (or may exist). How is that so?

When we state the categorial problem – what ontological category do musical works belong to? – it is arguable that the question, as it is stated, commits the complex question fallacy, since it presupposes that there is *one*, and only one, ontological category which subsumes *all* things that are musical works. But just as there is a working intuition telling us that there is ontological diversity within the class of *artworks* (namely, that some works are *singular*, identified with concrete physical objects, while some works are *multiple*, identified with no specific concrete physical objects), it is not clear that the class of musical works is ontologically homogeneous, even though our intuitions or *prima facie* beliefs about them point in that direction. It could be the case that some musical works are singular, even if most musical works are multiple (in the sense of having many actual or possible *instances*). Furthermore, even if in actuality there are no musical works which are singular entities, there is always the question of whether there *could* be such a thing. If there could be musical works that were singular entities, it would be merely contingent that none existed in the actual world and such contingent fact would not constrain our answer to the categorial problem.

To see this, consider a world were, contingently, there are no actual musical works (or no *instances* of musical works, nor acts of *indication* to which performances could be causally connected); suppose further that the class of musical works was indeed ontologically homogeneous, so that no musical work *could* be a singular entity. In such a world, it would still be *true* that musical works, should any exist or be instantiated, are ontologically homogeneous *and* multiple

(multiply instantiable) entities. Likewise, the world either contains musical works that are singular entities as one of its possibilities (one of its possible *ways of being*) or it doesn't. If it does, then the contingent non-existence of any such things makes no difference for the categorial problem: the class of musical works would still be correctly described as *not* ontologically homogeneous.

But what would be the relevance of *that* for our purposes? How would that have any bearing on the applicability of our Howellian-Searlean analysis of socially constructed kinds? My view is that it would not change the way the analysis applies to socially constructed items in general and to musical works in particular. This is due to the very ambiguity I have noted in the Searlean approach to the categorial problem: there is more than one ontological view compatible with the claim that we live in just *one* world; there is nothing in the idea of there being repeatable entities *as well as* non-repeatable ones to suggest that descriptions of those entities are descriptions of two distinct *worlds*. (We may think that the problem lies precisely in the vagueness of the term “world” in this context – is it an oblique way of speaking merely in *distinct kinds of entity*? Do we have different “worlds” if we postulate entities which are causally inert? – And it is not even clear whether abstract entities would be causally inert *if* they existed).

If the class of things which are musical works was ontologically diverse, in the sense of containing both singular and multiple entities, this would not affect the applicability of the formula *X* counts as *Y* in *C*, nor Howell's distinction between properties, patterns and types. At most, we would conclude that some musical works (the putative “singular” ones) would not be identifiable with types at all, though there would still be an *X* term (a pattern or combination of properties) which would receive a function *Y* (the musical work function) in a context *C*. Such items would still be individuated contextually, in virtue of the fact that a mere pattern or combination of properties would be insufficient to make it the case that there was a *musical work* – an entity which includes the pattern among its essential properties. And the class of such items would still constitute a more general type, namely, the type *singular* (non-multiply instantiable) *musical work*. Yet, a closer look at how a musical work which was also a singular entity would be like makes this possible distinction even less momentous: the best candidate we have for “singular musical work” are the kind of compositions which involve the direct manipulation of concrete sounds, such as a pure electronic montage which would have *playings* (or reproductions from a master recording) but not *performances* in the sense that Beethoven's Fifth Symphony has performances. Such works would be akin to *prints* made from a *matrix* (e.g. an etched plate) which have limited reproducibility (and this already challenges their classification as “singular”). No particular print is identifiable with *the work itself* and yet such works differ from most multiple artworks (poems, novels, musical works) in having their identity connected with a particular concrete artifact – the matrix – in the

sense that no print that is not the causal outcome of some action performed with the matrix (e.g. applying ink and running it through a printing press with paper) can be said to be a proper *instance* of the work. Prints from an etching are thus akin to *playings* or reproductions from a master tape, in the case of purely electronic compositions. Such works *are* multiple or repeatable, though they form a peculiar kind in including as essential the property of being causally connected to a concrete matrix or reproducing artifact. This strengthens the view that there could not be such things as purely singular (non-instantiable) musical works, in the sense that, say, paintings are thus conceived. At most, we would have two separate classes of multiple musical works – those that essentially involve a causal connection to a reproducing artifact or matrix (the master recording) and those that do not essentially involve such a connection (the “garden-variety” works for performance).

Whether the *X* term in “*X* counts as *Y* in *C*” is filled by a singular concrete item, a set of concrete items causally connected to a “matrix”, or an abstract entity of some sort (universal, kind, type, etc.), we have two invariable results from this “ontological scheme”: (a) the *X* term is always filled by some pattern or combination of properties (what varies is whether we focus on the pattern *qua* embodied in a singular object or in the pattern *qua* abstract object), and (b) the pattern or combination of properties is always *insufficient* to make it the case that there is an *artwork* or a *musical work*, over and above there being certain patterns (whether *qua* embodied in a singular concrete item, *qua* embodied in one or more concrete items causally connected to a matrix or reproducing artifact, or *qua* abstract entities).

What all this implies is that the result, for the categorial problem, of applying what I’ve been calling the Howell-Searle analysis is mostly *negative*: the analysis does not tell us what ontological category musical works belong to, but rather what category they do *not* belong or under which they cannot be subsumed *tout court*. The fact that Howell uses the term “type” to distinguish *something* from mere properties and patterns (combinations of properties) does not commit him (or anyone who employs his analysis) to a specific doctrine about the nature of “types”. His views are compatible, for instance, with Eddy Zemach’s theory of types as *concreta* (whether that theory is *true* is another story), as well as with several other theories of types. What it requires from us is that there is *something*, whether we call it “type” or some other name, which is distinct from mere properties and patterns and has certain essential properties in virtue of a connection with the directed states of social agents. Whatever that something is (concrete or abstract), it must be such that it can perform the functions expected of musical works and which (partially) *define* musical works as *musical* entities.

Whether or not the class of musical works is ontologically homogeneous, it does not seem to impinge on the way Howell's and Searle's analyses can be used to clarify the ontology of musical works *qua* socially constructed entities. Once we distinguish between properties, patterns and types, and use Searle's apparatus of causal-agentive functions and status-functions to clarify how a pattern *becomes* an artifact, we see that whether the patterns are thought of as "repeatable" abstract entities or *qua* embodied in a particular concrete object, if they are to constitute "works", they must fill the "X slot" in the "X counts as Y in C" formula. Of course, Searle himself insists that in his ontological scheme, which has been described as a "two-level ontology" – a "lower level" comprised of "brute facts" (those that exist independently of the directed states and coordinated representations of human agents) and a "higher level" of facts that depend on the imposition of functions, by human agents, on objects, events and other agents – *all facts* (including the fact that X counts as an artwork or as a musical work) must "bottom out" in brute facts, with which the former stand in a "continuous line":

"[There is a] continuous line that goes from molecules and mountains to screwdrivers, levers, and beautiful sunsets, and then to legislatures, money, and nation-states. The central span on the bridge from physics to society is collective intentionality, and the decisive movement on that bridge in the creation of social reality is the collective intentional imposition of function on entities that cannot perform these functions without that imposition."²⁷⁹

This prompts the question whether all "higher level" facts must bottom out in *concrete particular* entities, which would imply the exclusion of such ontological "abnormalities" as abstract entities (universals and types conceived as abstract), or whether abstract objects may be considered a legitimate "bedrock" where "higher level" facts could bottom out. This seems to depend more on results from general metaphysics than from the special ontology of musical entities. *If* the world is such that it contains abstract entities among its "basic constituents", then at least some facts about abstract entities would be as "brute" as any other such facts about concrete particulars (unless some conceptualist view of *abstracta* were true); and if the world is such that abstract entities cannot be counted among its basic constituents, then obviously there would be no facts (brute or otherwise) about abstract entities, except the fact that there would be no such things.

As we have seen, both the defender of a pure Platonist ontology (according to whom musical works are *identical* with *sound patterns*, abstract objects which are supposedly everlasting or

²⁷⁹ Searle (1995: p. 41).

timeless entities and discovered, not created by their composers) and the musical ontologist who is more mindful of descriptivist qualms (and thus more readily revisionist about traditional metaphysical categories) begin with *some* *prima facie* characterization of musical works. Pure Platonists stress the properties of hearability and repeatability, posing as musical ontology's main task that of explaining how musical works can have both properties – an explanation that would come from identifying the ontological category to which musical works belong. But to start with that list of two properties and to stress the metaphysically perplexing character of how they can belong to the same entity does not seem to have any legitimate priority over other ways of approaching the categorial problem. If we start with a list of three properties, say, that musical works, whatever else they may be, are a) hearable, b) repeatable and c) *creatable* (they are the *creations* of their composers), we merely add one more property which shares the same source as the previous two: our *prima facie* beliefs or intuitions about the nature of musical works. There seems to be no principled way of ascribing the property of creatability a lower metaphysical priority, in contrast with hearability and repeatability. The fact that it is *possible* that such intuition is mistaken (maybe musical works are, contrary to our intuitions, discoverable and not creatable) is of no consequence here, for we might equally be mistaken about the other two.

Maybe musical works are not, after all, repeatable entities – say, in virtue of the fact that *no* entities are, despite our intuitions – and maybe musical works are not, strictly speaking, hearable – one could argue that Wolterstorff's doctrine of “analogical predication”, endorsed by other Platonists such as Dodd, indeed amounts to a recognition that only *performances* of musical works, sound events, are hearable, and musical works cannot be identified with any particular sound event, the same way a chess game, a certain sequence of moves by players following the rules of chess, cannot be identified with the particular movements of particular pieces on a physical chessboard, which implies that no chess game can share certain properties with the physical objects with which it is played – e.g. no chess game is a *visible* entity.

And why would the purpose of explaining the perplexities involved in combining hearability and repeatability have some sort of metaphysical priority over explaining the perplexities involved in combining hearability, repeatability and creatability? If a nominalist ontology is correct, and musical works are concrete entities of some sort, there would be no perplexity involved in their being both hearable and creatable, while repeatability would be explained *away* rather than explained *simpliciter*. Metaphysical error is no more liable to occur in ascribing the third property than in ascribing the two properties emphasized by the pure Platonist. It seems that to ascribe priority to those two properties is already to be committed with a view on individuation (musical works, whatever they are, are individuated by combinations of *acoustic* properties), and to beg the

question against other ontological views with respect to the categorial problem. A principled way of downplaying creatability over hearability and repeatability or rather, of downplaying the combination of all three properties over the combination of hearability and repeatability, is that the pure Platonist needs to establish his approach to the categorial problem as the *default view*, but no such principled way seems to be available.

The thought that musical works are hearable and repeatable (that they are things which can have many instances) is no less *intuitive* than the thought that musical works are creatable (that they are things which are brought into existence by their composers). Since all such ascriptions of properties to musical works are liable to metaphysical error, none of the intuitions listed has a claim to special epistemic status. The belief that musical works are creatable has thus no less an equal claim to constrain our musical ontology than the belief that they are repeatable. Furthermore, the perplexities involved in explaining how hearability and repeatability can be combined in the same entity are no less perplexing nor epistemically special than the perplexities involved in explaining how creatability and repeatability can be combined in the same entity.

Here we can see that reliance on the Pragmatic Constraint (or some version of it thereof) is not the only way to bolster the epistemic status of creatability *vis-à-vis* hearability and repeatability. Analogy with other socially constructed entities – words and other symbols, languages, laws, games, etc. – is as much a legitimate source of epistemic strength as any Pragmatic Constraint would be. Intuitively, all socially constructed entities, because they are *constructed*, and because they depend on the existence of contingent entities, such as social agents and the social practices that provide the adequate context for the imposition of functions, are intuitively creatable entities, “tethered” to specific times and contexts. Intuitively, it makes no sense to speak of languages, laws, marriages and games in the absence of any social agents. The relation between social agents and socially constructed entities seems much stronger than the relation obtaining, say, between geometrical figures and concrete objects exemplifying them. The existence of triangles, their availability in “logical space”, so to speak, does not seem to depend on the existence of triangle-shaped objects, in the way that the existence of socially constructed entities seems to depend on the existence of social agents. Such analogy gives credence to the view that musical works are creatable, without generating a compromise with the truth of all or even most folk beliefs about musical works (the stronger view put forward by proponents of the Pragmatic Constraint).

There is a sense in which a language, a word or a game can be said to be “available in logical space”, independently of any social agents actually imposing functions on certain patterns of properties; being such that it contains a language like Latin or a game like chess are *ways the world can be*, possibilities that are open for any world where the existence of social agents is itself an

open possibility. But this is quite different from asserting that such entities *actually exist* in some way. By contrast, to say that triangles exist as abstract objects of some sort is to assert that triangles *exist* even if no concrete particular objects instantiating the triangular shape exist; their existence is thus conceived as no less *actual* than the existence of Mount Everest now. But it is one thing to say that the world *can* be such that it contains things like the game of chess or backgammon, and quite another to say that the world *is* such that it presently contains an entity which *is* a game, though it has never been instantiated in any historico-social context, via the actions of any social agents. The game of chess did not exist before it was *invented*, that is, before certain *patterns of actions* were encoded in constitutive rules, making such patterns of actions, or *playings*, actually available to social agents standing in appropriate relations (causal, historical) to the *originating context* (putatively, its invention by medieval Persians or Indians and further development into its present-day form) by which the game became actually available as a part of the “furniture of the world” or as a true denizen of social reality. It is not clear that modal truths must be explained by appeal to *possibilia* (sui generis objects), rather than some other theory that does not postulate such entities, and that *possibilia* are identical with abstract objects of some sort (actually existent non-concrete entities); modal truths might have a different explanation not involving the postulation of *abstracta*, a subclass of which would include socially constructed kinds, and which would make it the case that such kinds actually exist, floating free of social contexts and the directed states of social agents, awaiting instantiation in something like a Platonic heaven.

And what of the difference between causal-agentive artifacts and status-functional artifacts? Does *that* difference make a difference for our approach to the categorial problem?

We have seen that both kinds of artifact share the property of necessarily being the result of the imposition of functions by some agent on some item; they differ in virtue of the kinds of functional properties involved. Causal-agentive functions are performed in virtue of the physical structure and causal powers of the items that receive the function and do not require any coordinated beliefs, within a social group, concerning the *status* of such items *qua* tokens of the relevant artifactual kind. Status functions, by contrast, can only be performed in virtue of coordinated beliefs or shared representations within a social group. Causal-agentive functions must trace back to certain directed states of cognitive agents, namely, *insights* about dependence relations between certain base properties and functional properties that the base properties sustain (when the functional properties include *aesthetic functions*, we have the kinds of phenomena that fall under the causal-agentive concept of *art*). We have also seen that there is a subclass of artifacts that overlaps both kinds – these involve status-functions that *depend* or are *metaphysically subordinate* to previously existent

causal-agentive functions (e.g. the kinds of phenomena that fall under the status-functional concept of *art*). Both kinds of function share the property of depending on the directed states of cognitive agents.

But if causal-agentive artifacts are also a kind of social kind, how do they fit Howell's scheme and Searle's scheme simultaneously, since they do not depend on the (implicit or explicit) application of the formula "X counts as Y in C"?

The answer to that question lies in the interpretation of the verbal expression *counts as*: status-functional artifacts *count as* tokens of the relevant kind in virtue of shared beliefs or representations of social agents in the appropriate context (there are tokens of a given status-functional kind when people *behave as* such tokens exist, a behavior which is directed at some object, event or agent); causal-agentive artifacts *count as* tokens of the relevant kind in virtue of *being* tokens of the relevant kind, and the fact that they *are* tokens of the relevant kind is due to their standing in certain relations with the directed states of agents (namely, of standing in certain *causal* relations with *insights* about the sustaining of functional properties by base properties), independently of this fact obtaining in the absence of coordinated beliefs about the status of the item and independently of its obtaining cross-contextually (not being tied to a specific context or disjunction of contexts). Thus, for instance, a causal-agentive artwork consists in a pattern or combination of properties (whether or not we think of it as embodied in a physical object or event) – the *X* term – which *counts as* an artwork – the *Y* term – in a context *C* (which is derivatively or a fortiori true of its counting as such cross-contextually). One must be careful here not to confuse *counts as Y* – in the sense of *satisfying certain conditions for being Y* – with the *action of being counted as Y* by some agent or other.

To the extent that causal-agentive functions depend crucially on the physical structure and causal properties of the items embodying the relevant patterns or combinations of properties, it seems that they must always be imposed on particular concrete objects, which excludes putative entities such as abstract types or universals. Now, if this observation is correct, and if the class of musical works, like the class of artworks, overlaps both the class of causal-agentive artifacts and the class of status-functional artifacts (i.e. if there is both a causal-agentive concept of musical work and a status-functional concept of musical work, as I believe there is, with a partial overlapping of things that fall under one or the other), at least those things which are musical works in the causal-agentive sense essentially involve the imposition of functions (in this case aesthetic functions) on particular concrete items (in this case, sound events or patterns of acoustic properties). But this conflicts with our intuition that musical works cannot be identified with any given concrete particular sound event; in other words, it conflicts with our intuition that musical works are repeatable entities. However, this conflict can be dispelled once we realize that the imposition of aesthetic functions on

any given concrete item is subordinate to the imposition of those same functions on any *possible* concrete item which carries the relevant base properties (so it can sustain the aesthetic ones). For some artworks (e.g. works of painting), the first physical object that actually realizes the functions acquires a privileged status as *original* and any other object which replicates those aesthetic functions will be classified as a *copy*, reproduction or, in the worst case scenario, a *forgery*. But musical works can't be reproductively forged; their identity is not tied to a specific concrete sound event (e.g. the first performance of the work or even the first "inner hearing" of it by the composer). The imposition of musico-aesthetic functions operates on *types* of sound event, and it is by way of this more primordial operation that it can successfully operate on concrete particular sound events, which replicate the relevant pattern of acoustic properties. But the same is true of so-called "singular" artworks, like paintings and non-cast sculptures: a successful imposition of aesthetic functions on the first concrete item that will thereby acquire the status of *original* is subordinate to the imposition of those functions on any *possible* concrete item that is able to sustain the relevant aesthetic functions. There is a pattern of base properties on which aesthetic functions are imposed, and then a particular concrete item is generated which, in the best case scenario, will indeed sustain the envisioned aesthetic properties. The difference here between music and paintings (and other "singular" art forms) is that in the case of music no concrete particular sound event is given the status of *original*, and identified with *the* artwork itself.

All this resonates with Howell's distinction between properties, patterns and types: *all* artworks, whether they are singular or multiple, involve the imposition of functions on *patterns of* (base) *properties*; it is a matter of contingent social practice that for some art forms the first concrete item to instantiate the functional properties in virtue of instantiating the base properties is also considered the *original* and identified with the artwork *tout court*. This would seem to suggest that *all* artworks are really types, the distinction between singular and multiple art forms being rooted in a contingent social practice of identifying certain concrete items as the originals and thus with the artworks themselves. But there is a subtler and more promising approach that I will now describe.

Both artworks and musical works (at least those musical works which are also artworks), *qua* types in the Howellian sense, involve "an element of status-function", even if the explanation of why there is art and music in the first place is not status-functional through and through. This "element of status-function" would explain at least some of the Howellian features of types *qua* created and contextually individuated entities, distinct from mere patterns of properties.

Howell builds his view on initiated and indicated types drawing on Jerrold Levinson's theory of types (Levinson never uses the term "indicated type", but speaks rather of "indicated structures" and "initiated types"), his own attempt at reconciling the creatability intuition with the other two:

hearability and repeatability – the only ones the pure Platonist considers relevant for the categorial problem. An “indicated type” is a type which results from an intentional act of indication, which relates a composer, a musico-historical context and time, and a “sound structure”. The composer selects and arranges certain pre-existent sonic patterns in a certain order, and by thus selecting them endows them with significance, which includes both aesthetic significance and other kinds of “content” or “meaning” in a broad sense. This act of selection, ordering and endowment with significance is what indication basically consists in. Such a view contrasts with pure Platonism in that musical works are not identified with abstract sound structures, even though a sound structure (a pattern of sonic properties) is an essential component of a musical work. Rather, the work is the outcome of the contextualized actions of the composer, culminating in the indication of a pattern of sonic properties, a pattern that acquires, in virtue of that very act of indication, a significance or meanings (in a broad sense) it could not have by itself, nor across any variations of context, time, and composer. In fact, for Levinson, we would not get the *same* musical work if we made the sonic pattern fix while changing the musico-historical context, the time of composition and the composer. There have been several objections raised to Levinson’s approach (concerning how strictly musical works should be tied to times of composition and the view that musical works have times and composers as *constituents*, if we choose to take that language too literally) but I do not wish to dwell on them. The crucial idea I want to take from Levinson is the notion of *indication*, as it enters Howell’s conception of initiated types.

Indication of a sonic pattern, in this view, marks the *initiation* of a type – the musical work. Contrary to some objections that have been raised against the notion of indicated types, what results from the indication of a sonic pattern is not merely *the sonic pattern S as indicated by F* (the objection here is that this would give rise to a proliferation of “unwanted types”, e.g. *penicillin as discovered by Fleming*). Our interest in artworks generally and musical works in particular especially concerns their *significant* properties, what James O. Young (2001; 2014) has called their “cognitive content”. This cognitive content includes both their aesthetic significance and other artistically relevant properties. Upon its *indication*, the sonic pattern *acquires functions* it could not have on its own; these functions include not only its aesthetic functions (the function of having certain aesthetic properties *in virtue of* certain non-aesthetic properties *as envisioned in the composer’s insight*), i.e. the functions that make a sonic pattern into *music*, but also the function of making available (hearable, visible, readable, interpretable, etc.) other *significant* properties besides purely aesthetic ones. (Or maybe in a more careful formulation: besides purely *formal* aesthetic ones – since even sonic patterns not designed by human agents can have formal aesthetic properties). This is how I believe the products of *indication*, for things like artworks, should be

distinguished from ludicrous entities such as *penicillin as discovered by Fleming*, or *Mount Everest as pointed at by me*, which are often mentioned in order to rebut the Levinsonian notion of indicated type. The mere fact that Fleming discovered penicillin does not add to it any *functional properties* besides that of serving as an anti-bacterial agent or an antibiotic – but that is what penicillin already does independently of human agency; Fleming discovered a useful causal power of penicillin and this was put to use for human benefit. But the point of a composer’s indication of a sonic pattern is not merely to point out or put to use some previously existent causal powers of that pattern; the composer *endows* or *confers upon* the sonic pattern new functions: aesthetic, artistic and other sorts of meaningful functions. Even if two indistinguishable sonic patterns, one produced by a bizarre random natural event, the other envisioned by a human composer, share many of their formal aesthetic properties, the humanly designed one will have such properties... *by design*, and the fact that it has them by design will inevitably endow it with further properties, which will impinge on its total aesthetic profile in a variety of ways, all sensible to contextual and historical factors. It can determine other meaningful non-aesthetic properties which in their turn may have an effect on the total aesthetic profile of the work. For instance, that a tune was designed to be both elegant and humorous in a certain context may add to it meaningful properties; that it was designed by a certain individual in that context may add to its “meaning” in a way that makes it more aesthetically powerful than it would otherwise be. Compare this with the simpler fact that two superficially indistinguishable utterances by two individuals may differ in that one of them is ironic, in virtue of contextual factors.

Howell’s own illustration of this is Garbo’s utterance in *Grand Hotel*, “I want to be alone”, which exemplifies sentences that “are not only verbal patterns that carry expressive meanings and a historical content that go beyond their literal meanings”, but are also “tied to the particular individuals who first produced instances of them and got the practice going” (2002: pp. 118-119). This makes ulterior uses of such sentences “idiomatic” – their identity lies in their original performance by the utterer to whom further uses allude. Such sentences have functions (in virtue of which they *are* or *count as* sentences in the language) that are not tied to any specific individual utterer; but they also carry meanings (they acquire ulterior functions) that go beyond the former and are only intelligible in their connection with a particular context and the individual or individuals who uttered them in that context.

There is something similar going on with the imposition of functions on patterns of properties (including sonic patterns) to generate aesthetic and artistic functions, which go much beyond the purely formal aesthetic properties that those patterns could exemplify if taken as “indiscernible duplicates” produced by bizarre random natural events. Arthur Danto came close to this idea when

he said that artworks are essentially things that require some *interpretation: qua* objects of human *design*, instantiations of patterns are given *to be appreciated as products of an intentional operation on those patterns*, not merely as instantiations of the patterns *simpliciter*. Gregory Currie has made a similar point when he remarked:

“An interest in the aesthetics of artifacts is, for those cases where the distinction is a real one, an interest in something that unites both factors [activity and product]: an interest in the product-as-outcome-of-activity. That is why the aesthetic appreciation of nature *as genuinely natural* is so different from the aesthetic appreciation of art.”²⁸⁰

Artworks in general, and musical works in particular, have the function of generating this “interest in the product-as-outcome-of-activity”, a property which can be broken down in the several aesthetic, artistic and other meaningful functions the work acquires *qua* product of *indication*.

The practice of attaching meanings or significance to patterned objects/events thus differs considerably from the practice of simply making use of the previously existing causal powers of some given entity (as in the case of penicillin). When the Nazis adopted the hooked cross as their symbol, to use one of Howell’s examples once more (though not the most pleasant that comes to mind), they did not simply made use of some previously existing entity whose causal powers suited some of their purposes; they in fact *created* something (even though they did not create the geometric pattern of the hooked cross) by *imposing functions* on a pre-existent pattern. They made it henceforth possible for tokens of a certain geometrical pattern to *function* as conveyors of a specific meaning and thus introduced something in the world which did not previously exist. Artworks and musical works add further layers of complexity to this scheme, but are similar in that the functions added to the relevant patterns are not functions those patterns could perform *qua* disconnected from the directed states of human social agents.

We should now introduce a further distinction. All *indicated* types are, in virtue of that fact, also *initiated* types (they come into existence at a certain point in time); but not all initiated types are indicated types. The distinction lies in the fact that indicated types are tied to the actions of a specific individual or individuals (the contextualized actions that constitute *indication*). But some types are initiated, i.e. they are brought into being by social practices, without being tied to the actions of specific individuals in that manner. Some examples of these are words, languages, non-

²⁸⁰ Currie (2009: p. 18).

linguistic symbol-types, and other social and institutional kinds. The game of chess, for instance, is an initiated type, but its temporal initiation, its historical nature, is not tied to the actions of any specific individuals. No relevant property of the game of chess depends on our discovering that its initiation is tied to a given group or sequence of individual contributions rather than others. This is because the only “meaningful” properties of chess are tied to the constitutive rules that make *chess playings* possible, and these would not suffer any differences should we discover, say, that a certain variant of a rule was due to the contribution of the individual agent G rather than the individual agent F (compare with the case where a particular *chess move* is performed by a particular individual, in such a way that making *that* move, in *that* context, carries some further “meaning” – maybe an allusion to the *style* of another player, and similar things). The case is different (and substantially more complex) with artworks and musical works precisely in virtue of *meaning-functions* or *significance-functions* (more generally, properties that determine aesthetic, artistic or other forms of significance or “cognitive content”) that are liable to substantial variations given some variation in context or in the identity of the individuals who impose such functions on “patterns”.

We thus see that, even though the functional properties in virtue of which a given object or event may count (in the most primitive sense) as *art*, or a given sequence of sounds may count (in the most primitive sense) as *music* (before we complicate the scenario by introducing the concept of *work*) are causal-agentive (because aesthetic functions are causal-agentive), artworks and musical works perform a diversity of *meaning* or *significance* functions they could not perform exclusively in virtue of the intrinsic properties of the abstract patterns they “embody”; and this makes explicit the “element of status-function” present in them, and allows them to fit comfortably the Searlean formula “X counts as Y in C”, as well as making the Howellian distinction between patterns and types more intelligible.

Pure Platonists use a concept of *work* (in speaking of artworks or musical works) which is exclusively tied to the property of repeatability: any concrete sequence of sounds instantiates a pattern which can be instantiated by *other* sequences. If we think of entities such as Beethoven’s Fifth Symphony as mere sequences of sounds capable of multiple instantiation, the *work* concept will play no other role than that of signaling the fact that sequences are distinguishable from other sequences and can be “repeated” in the performances of musicians playing or singing from the score, from memory, or whatnot.

But the work concept plays a rather more substantial role than that. This fact may even lend some credence to the views of authors (e.g. Lydia Goehr) who hold that the work concept, or a *certain work concept* we express via the word “work”, is only operative for a specific period in the

history of music (however that period and that concept are to be precisely defined). This may be explained in terms of making more explicit the “elements of status-function” I alluded to before, that is, of precisely what functions we expect *indicated musical types* to perform in our culture. Plausibly, these expected functions will not be cross-temporally and cross-culturally invariable, and variations in the expected functions may determine the independently plausible result that some things we call “musical works” are *initiated* types without being *indicated* (these will be more like *words* and *games*) – we can think here of “works” in the oral tradition, whose identity is not tied to a specific individual or individuals and may even not have been “composed” in such a manner –, while others are, properly speaking, indicated types – those that fall under a more “thick” work concept, to employ Stephen Davies’ useful metaphor (though Davies was thinking specifically in terms of *sonic* properties when he introduced that metaphor, in how much “sonic detail” is specified by the composer and required of a performance in order for it to count as a *bona fide* instance of the “work”). For instance, works falling under the more “thick” concept may have the function of conveying to us something about the *style* of the composer, and how his individual style is itself related to a background of artistic activities (wider “styles” that span individual barriers and provide a background against which certain properties of a work are made visible or noticeable – think again of the case where a composer writes for a tuning system which does not coincide with the standardly adopted within his culture). Purely abstract acoustic “patterns”, by contrast, have no “styles”, and cannot, in and of themselves communicate anything whatsoever about some particular individual’s “situation” *vis-à-vis* his time, social and artistic context, personal history, his *oeuvre*, etc. In and of itself, no abstract sonic pattern “belongs” in a “tuning system”, since a given ordering of sounds can only be thought of as “belonging in a tuning system” relative to the directed states of human agents, to the way *we* organize sonic patterns *in order* for them to instantiate “meaningful” properties – even if this is just the property of forming certain kinds of intervals, scales and chords. None of these things – tuning systems, intervals, scales, chords – are pure acoustic patterns, but rather “intentional” entities: it is only at a higher order “level of description” (one involving the directed states of human agents) that an acoustic pattern counts as a scale or a chord. And once we move from the level of mere intervals, scales and chords to full blown musical “constructs” – melodies, harmonic progressions, etc. – the array of “meaningful properties” that such constructs may acquire *by human design* blows off the chart, so to speak.

It is important to note that though the pure Platonist about musical works is committed with the thesis that *all* musical works are individuated purely in terms of acoustic properties – this is the thesis of sonicism, in its “timbral” and “non-timbral” varieties – a contextualist about the ontology of music is not committed to the view that *all* musical works are individuated the same way, nor

does he have to deny that sonicism, timbral or otherwise, does in fact apply to the individuation of *some* musical works – maybe for the works of a certain period, or for works of a certain kind or style. The contextualist is not threatened by the existence of musical works for which the identity of the composer or other contextual factors are not identity-relevant. He is vindicated as long as it is true for *some* works that they require more “fine individuation” than what sonicism prescribes: that numerical identity for musical works is secured by sameness of normative acoustic properties in token sound events that purport to instantiate them.

Here is Julian Dodd’s own formulation of (timbral) sonicism:

“work *W* and work *W** are numerically identical if and only if they have exactly the same acoustic properties normative within them (where these acoustic properties include timbral properties). In other words, the condition that a sound sequence must meet to be a properly formed token of *W* is purely acoustic.”²⁸¹

Now, in light of what we have seen previously, this would be true if musical works were pure acoustic patterns, but that is no more plausible than to suppose that word-types are purely acoustic or graphic patterns; how could it be so, if such properties are insufficient to individuate words, if it is at least metaphysically possible that such properties be instantiated without there being any social agents making certain properties *normative* qua properties of a given kind?

The problem with this view is that it sees the *acts of making* (properties) *normative* as something purely *external* to the essential features of musical works qua musical works, somewhat like Jupiter’s moons are wholly external to the acts of discovery by which we come to have knowledge of them. This obscures the fact that to discover Jupiter’s moons is not to bring certain items into existence – say, Jupiter’s moons as “indicated” by Galileo – precisely because there is no essential normative component involved in the discovery of celestial bodies. To *indicate* a sonic pattern results in making certain acoustic properties *normative* for musical works, but this act of *making normative* is *essential* for musical works in a way it is not for the discovery of natural objects, such as moons and other celestial bodies. The way in which we count certain properties as “normative” for natural kinds differs relevantly from the way certain properties count as normative for socially constructed kinds. “Social construction” *consists* in ascribing normativity to certain properties, while the discovery of natural kinds consists in abstracting from nature certain preexisting “patterns” (e.g. being mostly made up of empty space and H₂O molecules) – patterns that have

²⁸¹ Dodd (2007: p. 240).

certain causal powers which are valuable to us, but that remain pretty much unaltered by the fact that we share beliefs about them (the fact that we *represent* the world as divided into natural kinds).

This view has the virtue (if it is a virtue) of not committing us to ontological monism about musical works (I take here “ontological monism” to mean the view that musical works must all belong to the same ontological category *and* that all musical works must be alike in terms of individuation, e.g., that *all* musical works must conform to the way that timbral sonicism conceives of individuation). Rejecting ontological monism about musical works, one may accept that timbral sonicism is true for *some* musical works but not all. This was in great measure already suggested by Stephen Davies (2001) and explained via the already mentioned metaphor of “thick” and “thin” works: there is no one true characterization of musical works for all periods and historical contexts, i.e. there is no *cross-historical* true ontological characterization of musical works, but rather a diversity of true characterizations, depending on *our interests*, which, in their turn, impose different *criteria of individuation* on musical works. This is not to introduce the ontological determination thesis of the social constructivist through the back door, though. It merely says that *what* types we are picking out by using the phrase “musical work” depends on what *musical work concept* is being deployed, and what concept is being deployed may vary with how “thick” our conception of *work* is, i.e., with how much properties we take as individuating of the “type” constituted by the sonic pattern (and normative variations of that pattern thereof) and the intentional operation of *indication* performed by human agents on patterns, endowing the resulting types with essential properties that the patterns cannot have by themselves.

Conclusion

We began our journey through the realm of musical metaphysics by exploring several approaches to the nature of artworks, taking note of the relevant connections, the similarities and asymmetries between theorization about the nature of artworks and the nature of music, as well as the link between these and conceptions artifacts and socially constructed kinds in general.

The purpose of doing so was to provide us with a way of bypassing the “ontological silent side-partner” of the “Pragmatic Constraint”, namely, social constructivism about socially constructed kinds – the view that the ontological nature of at least some socially constructed kinds is *determined* by the beliefs of “practitioners” or those who engage in the relevant practices that provide the entities under scrutiny a contextual bedrock.

I have tried at least to lend some plausibility to the idea that we can make meta-ontological realism compatible with the methodological constraint according to which our only sources of information concerning artworks generally and musical works in particular are the relevant practices: both “folk” and critical discourses on artworks and musical works. We can resist the thought that the beliefs which structure those practices are ontologically determinative, even if not *all* information we can gather from them can be found to be wrong, while maintaining that socially constructed things differ relevantly from natural kinds in being observer-dependent (or “intentionality-dependent”) phenomena.

Exploring the nature of socially constructed kinds, we have focused on the difference between social entities generally and *institutional* ones, which are a subclass of the former. John Searle’s theory of the construction of social reality was brought in as a device enabling us to make important distinctions, namely distinctions between *kinds of functional properties* determining different *artifactual categories* or different kinds of artifact. In doing so we arrived at a twofold categorization of artifacts: causal-agentive artifacts and status-functional artifacts (including overlaps between the two categories). The root of that categorization is the distinction between causal-agentive functions and status-functions. This distinction was proposed by Searle: causal-agentive functions are those an object has in virtue of an agent imposing that function on the object, which it performs in virtue of its physical structure and causal powers. Collective belief or representation is not required for these functions to be enacted. We have also used Nick Zangwill’s formulation of *aesthetic functionalism* about works of art, extending that analysis to artifacts generally, in order to clarify what makes an artifactual entity an entity *of its specific kind*: There are

certain insights an agent has, insights about relations of dependence between properties (though they need not be conceptualized in such explicit terms, or at all, by the agent), namely, *base physical properties* and *functional properties*; such that the former *sustain* the latter. All causal-agentive artifacts are defined (at least) by causal-agentive functions, which are properties that artifacts have in virtue of a causal connection to an agent's insights, as specified. This is (at least) part of what gives artifacts their "artifactuality".

But there are also other kinds of artifact. Kinds whose functions do not depend (or not only) on physical structure and causal powers but on *collective representations* about the very function the artifact is supposed to perform. Here the concept *artifact* is used in a broader sense, encompassing more than "human made concrete physical object", namely, *institutions* and putative abstract entities such as languages, words and other things with semantic or symbolic functions. One major example of this kind of artifact are *games*. Going against the grain of Wittgensteinian anti-essentialism, we have used Jim Stone's helpful attempt at providing a general definition of *game*, in terms of complex actions which include a *state of success*, a state which *counts* as success not because it is the "natural end" or purpose of the activity, but because *an arbitrary rule* establishes it as success. In other words, it is not success because we pursue it, rather we pursue it because there is a rule-governed practice of counting it has success. Games also entail *constitutive rules*: rules that do not merely regulate a previously existing activity but rather constitute the activity itself (chess rules make it the case that there are actions counting as chess moves). Thus games are institutional entities through and through. They conform to the Searlean formula for all institutional facts: *X counts as Y in C* (i.e. a certain object/event/agent *X* counts as having/performing the function *Y* in context *C*.) The analysis can be seen as an extension of Stone's definition of games for other institutional artifacts, analogously to the way Zangwill's aesthetic functionalism was expanded to provide a characterization of all causal-agentive artifacts. The fact that *X* counts as *Y* is not a consequence of *X being Y* in a rule-independent manner, but rather *X* counts as *Y* as a consequence of the rule determining that it does so *in the relevant context*. "Context" here can be specified as the conjunction of two further facts: the fact that applications of the rule both iterate indefinitely "upwards" and interlock "horizontally" with other institutional facts. Another way to put this point is to say that, typically, the relevant context for institutional facts includes a myriad other institutional facts.

Armed with these distinctions, we applied that expanded analysis to the institutional theory of art (broadly conceived). The result was that in order to be genuinely *institutional*, rather than just a "social-context" theory, a theory of art must conceive of art-activities in a way that makes them

relevantly similar to games and other institutional entities: there is a state of success (that of *X* counting as an artwork *Y*) which is not success in virtue of objects or events successfully performing certain causal-agentive functions, but rather because a conventionally arbitrary rule determines that the object/event acquires that status in the relevant context (e.g. *The snow shovel counts as art in the twentieth century artworld*). This raised doubts on whether artwork status for *all* artworks can be explained in this manner. Yet there seems to be a class of things called “artworks” that conforms well to the analysis. This is shown partly in Dickie’s reply to Wollheim’s Dilemma: there are all sorts of reasons to endow an object with institutional status, but what *makes* it have that status is not inherent to the reasons, but the fact that status was bestowed in the relevant context.

Such considerations lead us to consider the working hypothesis (defended by P. D. Magnus and Mag Uidhir but also anticipated by Zangwill) that “conceptual monism” about the art-concept (the view that there is only *one* true or useful concept of art) may be a philosophical dead end. Maybe there are radical differences within the class of artworks that call for distinct concepts (or conceptions) of art. Evidently, this begs the question of what makes several distinct artwork-*concepts* artwork-concepts, no matter how radical the differences they are supposed to cover. This in turn lead us to consider the following possibility: a twofold approach to the artwork-concept, in which both an “institutional” or *status-functional* concept of art and an “aesthetic” or *causal-agentive* concept of art coexist, but in such a way that the latter is *metaphysically and explanatorily prior* to the former. What this means is that status-functional art depends, for its existence and intelligibility, on the prior existence of causal-agentive art. Status-functional art cannot exist in a world entirely devoid of causal-agentive art. Also, status-functional art essentially involves the use of an explicit art-concept, while causal-agentive art does not: we count items of causal-agentive art as *art* because of our awareness that they perform the relevant functions *and* because we have a causal-agentive concept of art.

There are more causal-agentive functions, other than aesthetic functions, that artworks can have or perform. An expression theorist of art has a causal-agentive concept of art, for instance. Yet, many considerations and intuitions lead us to believe that although art can perform many functions, and that these may become important sources of artistic value, aesthetic functions mark a distinct boundary that separates artworks in the most primitive sense from more complex artistic objects that historically accumulate other functions as proper sources of value. Aesthetic functions also *explain* possession of other causal-agentive functions by artworks, such as the “clarifying” expression of emotions (Collingwood). Smashing a tea cup may be expressive of anger, but it is not art nor does it bear artistic value unless it does so in a way which explores the aesthetic dimension.

Representational functions add a layer of complexity, since they incorporate both causal-agentive and status-functional elements. It is the function of *representing* reality or *being about* something that makes possible the transition from causal-agentive art to status-functional art. Better yet, it is representational functions which allow art to become status-functional as well as causal-agentive. (For instance, a ready-made is unable to function as a ready-made without the ability to make a reference either to objects of the category of sculpture or to the category itself.) This adjusts well to the notion that aboutness is an essential feature of at least certain artworks, particularly those works of *avant-garde* art which are classified as *conceptual art*.

This analysis was furnished with a characterization of *aesthetic functions* as a subclass of causal-agentive functions. An aesthetic function is the function a given object or event has of *having* a certain aesthetic *property* – e.g. elegance, boldness, delicateness, gracefulness, daintiness, powerfulness, etc. – in virtue of certain non-aesthetic properties that *sustain* them. Here the “dependence base” of aesthetic properties can be thought of as *wide*, including non-perceptual contextual factors, or as *narrow*, including only “intrinsic” perceptual properties. There is no immediate pressure to choose one or other approach. Furthermore, the analysis is neutral among different theories of the *ontology* of aesthetic properties, which means that it is compatible with both “realist” and “anti-realist” approaches (whether we think that beauty and other aesthetic properties are “in the eye of the beholder” or that they are somehow part of the “furniture of the world” puts no apparent pressure on the analysis). It suffices for our purposes that aesthetic functions are causal-agentive, even if it turns out, as it is plausible, that imposing certain status-functions on an object (especially representational functions) may have an impact on its aesthetic profile. (Two objects with indistinguishable “formal” or appearance properties may differ aesthetically when they are differently considered as a representation and as a non-representational item.)

Having completed this journey through theories of art, institutional facts, games, artifacts and kinds of functional properties, the apparatus gathered is brought to bear on the nature of musical entities, in order to clarify what it is about musical entities that makes them *musical* rather than something else.

Drawing extensively on Roger Scruton’s well-known distinction between *sounds* and *tones* (in which *tone* becomes a technical philosophical concept with a restricted meaning) and his theory of *acousmatic listening* (a term that refers to aural experiences considered in abstraction from sound sources or from knowledge of sound sources), I’ve expounded a conception of what makes something *musical*, while avoiding a commitment with Scruton’s anti-realist inclinations. The result

is a view according to which the musical nature of a sound sequence is tied to aspects of human sound perception that are not exclusive of aural experiences but constitute a universal feature of how human beings experience nature and is an essential aspect of how natural environments also become a part of our *cultural* environment. This is the ability to explore *cross-domain mappings*, an ability that James O. Young has characterized as an ability to enact a “transfer of concepts derived from one sensory modality to experiences derived from another sensory modality” (Young: 2014, p. 19). It is in virtue of this ability that we can have aural experiences that we describe in terms of height, motion and emotion, when nothing literally *moves*, is *higher or lower* or *emotes* in the way those descriptions suggest, should they be taken as accurate literal descriptions of what is actually going on in the sounds.

The view I suggest is that there is a kind of *basic musical experience*, that is, a kind of aural experience involving *basic musical features* – these are the features of rhythm, melody, harmony and, arguably, timbre (though the inclusion of timbre here is not without its problems) – which is the core of any more complex forms of musical experience. Scruton famously describes the experience of rhythm, melody and harmony in terms of *spatial metaphors*, arguing for the more contentious view that descriptions of music at this “fundamental” level are necessarily metaphorical (that no predicates of height, motion and emotion apply *literally* to music). This thesis has had its detractors, but nothing of great weight hangs on our accepting or rejecting the necessity of metaphorical description for basic musical experiences. Whether or not what Zangwill has termed “the essential metaphor thesis” is true, or whether descriptions of music “at the fundamental level” are “semantically extended” though literal (which is a point about the *language used*, not the things described), the *experiences* to which those descriptions apply involve the mentioned perceptual of cross-domain mappings or “conceptual transfers” between sense modalities, independently of whether the *predicates* we use to describe those experiences can appropriately be said to be metaphorical or merely “semantically extended” literal terms. Though I am sympathetic to the essential metaphor thesis, I somehow downplay the importance of this point about language. I discussed the views of Scruton and Peacocke, according to whom “metaphorical” applies *to the perceptual experiences themselves*, not just to the language used to describe those experiences. This may be just a confused way of speaking about the mechanism of cross-domain mapping, and probably we’ll do best by reserving the term “metaphor” for linguistic phenomena, for the sake of clarity. Another option would be to use the concept of *aesthetic conceptual transfer* as a neutral placeholder to describe the phenomenon of “aesthetic metaphor”.

The remaining part of my proposal is that the experiences of rhythm, melody and harmony are themselves dependent on this mechanism of cross-domain mapping as applied to aural experiences, a phenomenon that lends itself to be described in Scruton's language of "metaphorical transfer" or Peacocke's "perceiving metaphorically-as", even if we reject the "essential metaphor thesis". In fact, a resemblance theory of music's emotional expressiveness is as amenable as the "essential metaphor thesis" to incorporate the notion of a cross-domain mapping mechanism as part of the explanation of why we experience *resemblances* between music and non-musical phenomena (Schroeder). But the truth of a resemblance theory would provide a reason not to consider the predicates of height, motion and emotion as metaphorical.

I was particularly interested in is the way in which this mechanism of "conceptual transfer" or "cross-domain mapping" can *determine* the aesthetic properties of sound events. All kinds of sound events will have aesthetic properties or qualities, but my suggestion is that the mechanism of "conceptual transfer" responsible for our hearing *movement* and other *spatial relations* in sounds (or sonic properties most naturally described by us in motion and spatial terms) determines a *subclass* of sonic aesthetic properties – all the aesthetic properties of a sound event whose dependence base includes those properties a sound event has in virtue of being perceived under the mechanism of cross-domain mapping (in less words: everything else we hear in sounds in virtue of hearing "movement" and "spatial relations" in it) – which in their turn define what it is for a sound event to be a musical sound event. We can either use the concept *music* to pick out that peculiar kind of aural experience, or we can use the concept in a much wider sense. Whichever coincides with the folk concept is just a matter of contingent cultural practice. The point is that there *is* that peculiar kind of aural experience, however we choose to refer to it.

And why is it important to nail down that peculiar kind of aural experience? Because the intentional design of sound events generating such experiences will be a special case of imposing *causal-agentive aesthetic functions* on those sound events. An agent has the relevant insights, that certain sonic aesthetic properties will be sustained by certain non-aesthetic properties (sounds produced with such and such instruments, in such and such a manner) and the fact that the sound event is causally connected with those insights, making it the case that the sound events produced under the guidance of such insights have the aesthetic properties they have *by design*, also makes it the case that the sound events are *music* in the *causal-agentive* sense of the concept (replicating here the same twofold conceptual structure we have suggested for the *art-concept*). Thus, the *primitive* (causal-agentive) sense in which something is *music* is 1) that of exhibiting the kind of aesthetic properties whose dependence base includes the "cross-domain" perception of movement and spatial

relations *in sounds* and 2) that of having such properties *by design*, tracing back to an agent's insights concerning dependencies between aesthetic and non-aesthetic properties of sounds.

Once this causal-agentive concept of music is in place, the stage is set for a status-functional concept (analogously to the status-functional concept of art) to eventually emerge. As with the status-functional concept of art, the status-functional concept of music will crucially involve the notion of *aboutness* – sound events not possessing the properties definitive of the causal-agentive concept of music are endowed with the status of music and thus made to be heard against a *background expectation* of hearing tonal relationships which are absent. The twofold structure of the concept of music can be construed as a disjunctive definition, in the style of that presented by Andrew Kania, which is here object of a survey, together with other two approaches on the definition of music – that of Jerrold Levinson and Andy Hamilton.

The next phase of our exploration is to draw some consequences, for the problem of the *ontology* of musical entities, from the above – combining aesthetic functionalism and the Searlean analysis of social construction to produce both a twofold art-concept and a twofold music-concept. While the problem of definition focuses on what properties something must have for it to be music, the problem of ontology focuses on what kinds of things musical entities are (also referred to as “the problem of the nature of musical works”, assuming here that musical *works* are the only kind of musical entities around), and what their *conditions of identity* are; e.g. what makes any two distinct sound events *S* and *S** be *bona fide* instances of one and the same work *W*.

One major focus of discussion is whether musical works are to be identified with abstract entities of some sort. Realist views about abstract entities see the positing of such entities as a crucial step in explaining a metaphysically perplexing feature that some items seem to exemplify, namely, repeatability. Some musical ontologists (e.g. Kivy, Dodd) consider repeatability to be the most metaphysically interesting and perplexing feature of musical works, so that the main task of musical ontology would be to explain how musical works can be both audible and repeatable. Given that concrete particular entities are by definition non-repeatable (they persist in time but do not “recur” in space, i.e., they only occupy one region of space at a time), if musical works are indeed repeatable, they must be abstract. But hearability seems to be a property that pertains only to concreta, since no object which has no spatial parts (and abstracta are typically conceived as non-spatial entities) can causally interact with the environment in order to produce audible sounds. Each performance of a musical work is a compound of unique, unrepeatable sounds (*soundings* are concrete particular events); but *in each performance* we claim to hear *the work*, so the work is what *repeats* (or gets repeated) *in* the concrete sound events that count as performances of the work. So

what exactly is the musical *work* (that which recurs in every *bona fide* performance)? This is the “categorical problem” of musical ontology: what *ontological category* should we place musical works in? The further problem of what makes an instance or performance of a work a *bona fide* one is what we call the “individuation problem”: what conditions must a sound event *S* satisfy in order for it to be a genuine instance of work *W* rather than work *W*^{*}?

There are many ontological theories of musical works on offer: some conceive musical works as *abstract types*: entities akin to *universals*, except that, like concrete particular entities they are objects of predication while themselves are not predicated of other entities; some conceive them as concrete entities, sets of concrete entities or mereological fusions of concrete entities; some conceive them as *action-types* (in which case they are also “abstract”) or as action-tokens (in which case they are concrete); others yet conceive them as *mental entities* of some sort (e.g. “ideas” in the minds of composers or listeners). There are theories that eliminate musical works altogether as denizens of a sound ontology (a “categorical map” of what there is); others yet conceive of musical work discourse as *fictional* (a kind of discourse which does not aim at truth but is nevertheless useful in some way). This abundance of views on offer is referred to by some (e.g. Thomasson) as an “embarrassment of riches”. I do not share the view that the riches are embarrassing. It is natural that in a period where people didn’t pay much attention to musical ontology, the hypotheses on offer were scarce. This does not mean that the discipline was more epistemically virtuous back then, or that it will become more epistemically virtuous once we eliminate or explain away most of the alternatives (we can never be assured we have done so for the right reasons). There is something to be learned in finding the pros and cons of each rival theory, something that goes well beyond the mere fact of one of them, or none, being right or true.

Focusing on the connections and interactions between the problems of definition and ontology, I try to determine whether the concept of *work* carries some ontological constraints, other than referring to a (presumably) repeatable and hearable entity. The question is then, how would we situate *works* (of music) in our previous categorization of artifacts? Given the twofold structure of the music-concept we sketched earlier, we see that *functional properties* are always included in the defining properties of music, whether these are causal-agentive or status-functional. Functional properties fit what Robert Howell has called “impure properties” (i.e. properties that “essentially involve” contingent particulars). In this case, functional properties make essential reference to human (or other) agents cognitively capable of imposing functions on objects – there are no functional properties unless there are either 1) insights about relations of dependence between properties, causally upstream of certain objects that exemplify the “supervenient” functional

properties in virtue of exemplifying the base properties which sustain the former, or 2) shared representations concerning the *status* *Y* of a certain object *X* in a context *C*. Defenders of the creatability of types (like Howell) appeal to impure properties as examples of non-eternal or temporally initiated properties (properties that cannot exist until the contingent entities which they “essentially involve” begin to exist). The idea is that if at least some types include functional properties as essential properties, and if functional properties are indeed “impure”, then the types specified by functional properties are temporally initiated. Such a result would secure the creatability of at least some types and would break the deadlock between defenders of “discoverism” about musical works (the view that musical works are *discovered eternal sound structures* rather than *creations* of their composers), a view typically endorsed by Platonists.

Defenders of the eternity of all types (such as Dodd) object to this, insisting that all properties, including the “impure” ones, are eternal. On Dodd’s account, a “property associate” of a type *K* is merely a condition something has to satisfy in order for it to be a token of *K*. The *actual* satisfaction of the condition is not itself a condition on the existence of the type, since types, on Dodd’s view, do not have their tokens or instances essentially. Thus, the “impure” property of *being a son of Lincoln* would exist, only uninstantiated (and uninstantiable) in Ancient Greece. The same would go for any types defined by such properties.

I have objected that this makes the implausible demand that modal truths be explained by appeal to abstract entities. To take an example, the impossibility of there being round squares does not seem (at least not obviously) to require the actual existence of a (non-instantiable) property of *being a round square*. Going back to works of music, the mere modal truth that the world is such that it *can* contain agents capable of imposing functions on objects, being true even when there are no human agents in the world, does not imply the truth of the view that there *actually exist* all musical works (and other socially constructed entities) that *can* be created (or “indicated”) by human agents once the world contains such agents. Making an analogy with symbol-types, it is plain that once the world contains human agents using geometrical patterns to perform semantic or representational functions, the possibility that *any* preexisting geometrical pattern be used as some form of symbol is included among the “modal structure” of the world, among the possibilities that are open within the world. But this is not the same as saying those entities *actually exist*. Something new is brought into the world when human agents actually impose functions on preexisting “patterns” of properties, be they geometrical, acoustic or otherwise. Using Howell’s distinction between types, patterns and properties, and combining it with Searle’s account of social construction, we’ve sketched an amplified view of the “indication” of “structures” (or patterns), one that Jerrold Levinson had

already used in his attempt to justify the intuition that musical works are created rather than the discovered, eternally existing entities the pure Platonist imagines.

Another point was that the concept of *work* brings with it “an element of status-function”, such that whatever functions we *expect* works to fulfill may have a bearing on what ontological characterization is appropriate to them. This would make the nature of musical works sensible to contextual factors in a way that *monist* ontologies of music have difficulty in accounting for. In other words, there is arguably a *false dichotomy* in the alternative that either the identity of composers is or is not (for all times and contexts) relevant for musical work identity. Coupling Howell’s analysis with Searle’s, and understanding the *indication* of preexisting patterns, resulting in the existence of “initiated types”, in terms of the imposition of a *status* (Y) on a pattern of properties (X) in a relevant context (C), we see that variability in the “content” of the Y term will have consequences at the level of the ontological description of the resulting types (musical works). Maybe some works, more “thin” in institutional functions, are appropriately accounted for by Julian Dodd’s “timbral sonicism” (the thesis that only timbrally inflected acoustic properties are identity-relevant for musical works), while others, more “thick” in institutional functions, are not, and require more “fine individuation”. Whatever else types may be, whether or not they are “abstract”, those types that are *works* are indeed temporally initiated, for it takes more than the mere fact that the world is such that it *can* contain those entities for it to be the case that such entities actually exist. Another point we focused on is that whatever functions we *expect* works to fulfill may have a bearing on what ontological characterization is appropriate to them. This would make the nature of musical works sensible to contextual factors in a way that *monist* ontologies of music have difficulty in accounting for.

Nothing in the course of our journey made us lean more forcibly towards a conception of musical works as abstracta or as concreta. The language of “types” is not itself decisive, since there are views on offer that conceive of types as concreta (e.g. Eddy Zemach’s eccentric ontology of types), and even if one such view is inadequate, there is no guarantee that it cannot be strengthened or refined. But the question whether a conception of types as non-abstract is true is not a *musical-ontological* question, except to the extent that something peculiar about *music* could make us lean toward such a conception.

What I hope came out more clearly is that musical works can’t be identified with *pure abstract acoustic patterns* (or decontextualized combinations of acoustic properties), independently of the directed states (insights, beliefs, experiences) of human agents. Both what makes any musical entity *music* and what makes a musical work be a musical *work* are functional properties, properties

essentially involving the existence, actions, insights and experiences of human agents. For both art and musical objects, to be a *work* is to *function* as such, i.e., it is to perform certain roles in the “culture” (the relevant “practice-saturated” context); and whatever roles are expected from something fulfilling the work-function in a given culture (what roles will specify a “thinner” or “thicker” work-function) is something prone to change or vary across times and contexts. And this means that ontological pluralism, rather than any monistic account of music’s ontology, is the more promising approach to take on the subject.

Bibliography

- Abell, C. (2005) "Against Depictive Conventionalism", *American Philosophical Quarterly*, Vol. 42, N° 3, pp. 185-197.
- Abell, C. (2011) "Art: What it Is and Why it Matters", *Philosophy and Phenomenological Research*, Vol. 85, N° 3, pp. 671-691.
- Alperson, P. (ed.) (1987) *What Is Music: An Introduction to the Philosophy of Music*, Penn State University Press.
- Ball, P. (2010) *The Music Instinct: How Music Works and Why We Can't Do Without It*, Oxford University Press.
- Beardsley, M. (1982) *The Aesthetic Point of View*, Cornell University Press.
- Bell, C. (1914) *Art*, New York, Capricorn Books, 1958.
- Budd, M. (2003) "Musical Movement and Aesthetic Metaphors", *British Journal of Aesthetics*, Vol. 43, N° 3, pp. 209-223.
- Burke, E. (1990) *A Philosophical Enquiry into the Origin of our Ideas of the Sublime and Beautiful*, Oxford University Press.
- Calvino, I. (2010) *The Invisible Cities*, trans. William Weaver, Random House.
- Carroll, N. (1993) "Historical Narratives and the Philosophy of Art", *Journal of Aesthetics and Art Criticism*, Vol. 51, N° 3, pp. 313-326.
- Copland, A. (2011) *What to Listen for in Music*, Penguin Group.
- Currie, G. (1989) *An Ontology of Art*, St Martin's Press.
- Currie, G. (2009) "Art for Art's Sake in the Old Stone Age", *Postgraduate Journal of Aesthetics*, Vol. 6, N° 1, pp.
- Danto, A. (1964) "The Artworld", *The Journal of Philosophy*, Vol. 61, N° 19, pp. 571-584.
- Danto, A. (1973) "Artworks and Real Things", *Theoria*, Vol. 39, N° 1-3, pp. 1-17.
- Danto, A. (1981) *The Transfiguration of the Commonplace*, Harvard University Press.

- Davidson, D. (1978) "What Metaphors Mean", *Critical Inquiry*, Vol. 5, N° 1, Special Issue on Metaphor, pp. 31-47.
- Davies, D. (2004) *Art as Performance*, John Wiley & Sons.
- Davies, S. (1988) "A Defence of the Institutional Theory of Art", *The Southern Journal of Philosophy*, Vol. 26, N° 3, pp. 307-324.
- Davies, S. (1990) "Functional and Procedural Definitions of Art", *Journal of Aesthetic Education*, Vol. 24, N° 2, pp. 99-106.
- Davies, S. (1991) *Definitions of Art*, Cornell University Press.
- Davies, S. (1997) "First Art and Art's Definition", *The Southern Journal of Philosophy*, Vol. 35, N° 1, pp. 19-34.
- Davies, S. (2001) *Musical Works and Performances: A Philosophical Exploration*, Clarendon Press.
- Davies, S. (2004) "The Cluster Theory of Art", *British Journal of Aesthetics*, Vol. 44, N° 3, pp. 297-300.
- Davies, S. (2008) "Musical Works and Orchestral Color", *British Journal of Aesthetics*, Vol. 48, N° 4, pp. 363-375.
- Davies, S. (2010) *Philosophical Perspectives on Art*, Oxford University Press.
- Davies, S. (2011) *Musical Understandings And Other Essays in the Philosophy of Music*, Oxford University Press.
- Davies, S. (2012a) "On Defining Music", *The Monist*, Vol. 95, N° 4, pp. 535-555.
- Davies, S. (2012b) *The Artful Species: Aesthetics, Art and Evolution*, Oxford University Press.
- Davies, S. et al (eds.) (2009) *A Companion to Aesthetics* (2nd edition), John Wiley & Sons.
- De Clercq, R. (2008) "The Structure of Aesthetic Properties", *Philosophy Compass*, Vol. 3, N° 5, pp. 894-909.
- Devitt, M. & Sterelny, K. (1999) *Language and Reality: An Introduction to the Philosophy of Language*, MIT Press.

- Dickie, G. (1964) "The Myth of the Aesthetic Attitude", *American Philosophical Quarterly*, Vol. 1, N° 1, pp. 56-65.
- Dickie, G. (1974) *Art and the Aesthetic: An Institutional Analysis*, Cornell University Press.
- Dickie, G. (1984) *The Art Circle: A Theory of Art*, New York, Haven Publications.
- Dickie, G. (1997) "Art: Function or Procedure: Nature or Culture?", *The Journal of Aesthetics and Art Criticism*, Vol. 55, N° 1, pp. 19-28.
- Dickie, G. (1998) "Wollheim's Dilemma", *British Journal of Aesthetics*, Vol. 38, N° 2, pp. 127-135.
- Dilworth, J. (2008) "The Abstractness of Artworks and Its Implications for Aesthetics", *The Journal of Aesthetics and Art Criticism*, Vol. 66, N° 4, pp. 341-353.
- Dodd, J. (2000) "Musical Works as Eternal Types", *British Journal of Aesthetics*, Vol. 40, N° 4, pp. 424-440.
- Dodd, J. (2002) "Defending Musical Platonism", *British Journal of Aesthetics*, Vol. 42, N° 4, pp. 380-402.
- Dodd, J. (2007) *Works of Music: An Essay in Ontology*, Oxford University Press.
- Dodd, J. (2008) "Musical Works Ontology and Metaontology", *Philosophy Compass*, Vol. 3, N° 6, pp. 1113-1134.
- Dodd, J. (2010) "Confessions of an Unrepentant Timbral Sonicist", *British Journal of Aesthetics*, Vol. 50, N° 1, pp. 34-52.
- Dodd, J. (2012) "Adventures in the Metaontology of Art: Local Descriptivism, Artifacts and Dreamcatchers", *Philosophical Studies*, Vol. 165, N° 3, pp. 1047-1068.
- Dutton, D. (2006) "A Naturalist Definition of Art", *The Journal of Aesthetics and Art Criticism*, Vol. 63, N° 3, pp. 367-377.
- Dutton, D. (2009) *The Art Instinct: Beauty, Pleasure and Human Evolution*, Oxford University Press.
- Evans, G. & Altham, J. E. J. (1973) "The Causal Theory of Names", *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 47, pp. 187-225.

- Gaut, B. (2005) "The Cluster Account of Art Defended", *British Journal of Aesthetics*, Vol. 45, N° 3, pp. 273-288.
- Goehr, L. (1989), "Being True to the Work", *The Journal of Aesthetics and Art Criticism*, Vol. 47, N° 1, pp. 55-67.
- Goehr, L. (2007) *The Imaginary Museum of Musical Works: An Essay in the Philosophy of Music*, Oxford University Press.
- Goodall, H. (2013) *The Story of Music*, Random House.
- Goodman, N. (1976) *Languages of Art: An Approach to a Theory of Symbols*, Hackett Publishing.
- Gracyk, T. (2013) *On Music*, Routledge.
- Hamilton, A. (2007a) *Aesthetics and Music*, Continuum International Publishing Group.
- Hamilton, A. (2007b) "Music and the Aural Arts", *British Journal of Aesthetics*, Vol. 47, N° 1, pp. 46-63.
- Hanslick, E. (1986) *On the Musically Beautiful* (trad. Geoffrey Payzant), Hackett.
- Howell, R. (2002) "Types, Indicated and Initiated", *British Journal of Aesthetics*, Vol. 42, N°2, pp. 105-127.
- Iseminger, G. (2004) *The Aesthetic Function of Art*, Cornell University Press.
- Kania, A. (2008) "The Methodology of Musical Ontology", *British Journal of Aesthetics*, Vol. 48, N° 4, pp. 426-444.
- Kania, A. (2010) "Silent Music", *The Journal of Aesthetics and Art Criticism*, Vol. 68, N° 4, pp. 343-353.
- Kania, A. "Definition", in Kania & Gracyk (eds.) (2011).
- Kant, I. (2007) *Critique of Judgement*, trans. Meredith, Oxford University Press.
- Kingsbury, J. & Mckeown-Green, J. (2009) "Definitions: Does Disjunction Mean Dysfunction?", *Journal of Philosophy*, Vol. 106, N° 10, pp. 568-585.
- Kivy, P. (1975) "What Makes 'Aesthetic' Terms Aesthetic?", *Philosophy and Phenomenological Research*, Vol. 36, N° 2, pp. 197-211.

- Kivy, P. (1993) *The Fine Art of Repetition: Essays in the Philosophy of Music*, Cambridge University Press.
- Kivy, P. (ed.) (2004) *The Blackwell Guide to Aesthetics*, John Wiley & Sons.
- Kristeller, P. O. (1965) "The Modern System of the Arts", in *Renaissance Thought II: Papers on Humanism and the Arts*, P. O. Kristeller (ed.), New York, Harper & Row.
- Lamarque, P. (2010) *Work and Object: Explorations in the Metaphysics of Art*, Oxford University Press.
- LePoidevin et al (eds.) (2009) *The Routledge Companion to Metaphysics*, Taylor & Francis Group.
- Levinson, J. (1980) "What Musical Work Is", *The Journal of Philosophy*, Vol. 77, N° 1, pp. 5-28.
- Levinson, J. (1990) *Music, Art and Metaphysics: Essays in Philosophical Aesthetics*, Oxford University Press.
- Levitin, D. (2010) *The World in Six Songs: How the Musical Brain Created Human Nature*, Aurum Press.
- Lewis, D. (1969) *Convention: A Philosophical Study*, John Wiley & Sons.
- Mag Uidhir, C. & Magnus, P. D. (2011) "Art Concept Pluralism", *Metaphylosophy*, Vol. 42, N° 1-2, pp. 83-97.
- Mag Uidhir, C. (ed.) (2013) *Art and Abstract Objects*, Oxford University Press.
- Mandelbaum, M. (1965) "Family Resemblances and Generalization Concerning the Arts", *American Philosophical Quarterly*, Vol. 2, N° 3, pp. 219-228.
- Margolis, E. & Laurence, D. (eds.) (2007) *Creations of the Mind: theories of artifacts and their representation*, Oxford University Press.
- Marmor, A. (1996) "On Convention", *Synthese*, Vol. 107, N° 3, pp. 349-371.
- Mathiessen, T. J. (1999) *Apollo's Lyre: Greek Music and Music Theory in Antiquity and the Middle Ages*, University of Nebraska Press.
- McGinn, C. (1997) *Ethics Evil and Fiction*, Oxford, Clarendon Press.
- Miller, R. (1992) "A Purely Causal Solution to One of the Qua Problems", *Australasian Journal of Philosophy*, 70, pp. 425-434.

- Millikan, R. G. (1998) "Language Conventions Made Simple", *The Journal of Philosophy*, Vol. 95, N° 4, pp. 161-180.
- Mithen, S. (2005) *The Singing Neanderthals: The Origins of Music, Language, Mind and Body*, Harvard University Press.
- Moravcsik, J. "Why Philosophy of Art in Cross-Cultural Perspective?", *The Journal of Aesthetics and Art Criticism*, Vol. 51, N°3, pp. 425-435.
- Nettl, B. (2005) *The Study of Ethnomusicology: Thirty One Issues and Concepts*, University of Illinois Press.
- Peacocke, C. (2009) "The Perception of Music: Sources of Significance", *British Journal of Aesthetics*, Vol. 49, N° 3, pp. 257-275.
- Pettersson, A. (2009) "P. F. Strawson and Stephen Davies on the Ontology of Art. A Critical Discussion", *Organon F16*, n°4, pp. 615-631.
- Putnam, H. (1992) *Realism With a Human Face*, Harvard University Press.
- Rohrbaugh, G. (2003) "Artworks as Historical Individuals", *European Journal of Philosophy*, Vol. 11, N° 2, pp. 177-205.
- Rosar, W. (2002) "Film Music – What's in a Name?", *The Journal of Film Music*, Vol. 1, N° 1, pp. 1-18.
- Rudner, R. (1950) "The Ontological Status of the Esthetic Object", *Philosophy and Phenomenological Research*, Vol. 10, N° 3, pp. 380-388.
- Sacks, O. (2007) *Musicophilia: Tales of Music and the Brain*, Knopf.
- Sainsbury, R. M. (2009) *Fiction and Fictionalism*, Routledge.
- Sainsbury & Tye (2012) *Seven Puzzles of Thought And How to Solve Them: An Originalist Theory of Concepts*, Oxford University Press.
- Sarah Hegenbarth and Vid Simonti (2011) "Beauty and Meaning; Music and Morality: An Interview With Roger Scruton", *Postgraduate Journal of Aesthetics*, Vol. 8, N° 1.
- Sartwell, C. (2004) *Six Names of Beauty*, Routledge Chapman & Hall.
- Schroeder, S. (2013) "Music and Metaphor", *British Journal of Aesthetics*, Vol. 53, N° 1, pp. 1-19.

- Scruton, R. (1983) *Aesthetic Understanding: Essays in the Philosophy of Art and Culture*, Methuen.
- Scruton, R. (1997) *The Aesthetics of Music*, Oxford University Press.
- Scruton, R. (2004) "Musical Movement: A Reply to Budd", *British Journal of Aesthetics*, Vol. 44, N° 2, pp. 184-187.
- Scruton, R. (2007b) "In Search of the Aesthetic", *British Journal of Aesthetics*, Vol. 47, N° 3, pp. 232-250.
- Scruton, R. (2009) *Understanding Music: Philosophy and Interpretation*, Continuum.
- Scruton, R. "Rhythm, Melody and Harmony", in Kania & Gracyk (eds.) (2011).
- Scruton, R. (2012) *Modern Philosophy: An Introduction and Survey*, A & C Black.
- Searle, J. (1995) *The Construction of Social Reality*, Simon and Schuster
- Searle, J. (2010) *Making the Social World: The Structure of Human Civilization*, Oxford University Press.
- Sibley, F. (1959) "Aesthetic Concepts", *The Philosophical Review*, Vol. 68, N° 4, pp. 421-450.
- Sibley, F. & Benson, J. (2001) *Approach to Aesthetics: Collected Papers on Philosophical Aesthetics*, Clarendon Press.
- Stachelhaus, H. (1991) *Joseph Beuys*, Perseus Distribution Services.
- Stecker, R. (1997) *Artworks: Definition, Meaning, Value*, Penn State Press.
- Stock, K. & Thomson-Jones, K. (2008) *New Waves in Aesthetics*, Palgrave Macmillan.
- Stone, J. (1994) "Games and Family Resemblances", *Philosophical Investigations*, Vol. 17, N° 2, pp. 435-443.
- Suits, B. (2014) *The Grasshopper: Games, Life and Utopia*, Broadview Press. (3rd ed.)
- Tanner, M. & Budd, M. (1985) "Understanding Music", *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 59, pp. 215-248.
- Tatarkiewicz, V. (1971) "What is Art? The Problem of Definition Today", *British Journal of Aesthetics*, Vol. 11, N° 2, pp. 134-153.
- Tatarkiewicz, V. (2006) *History of Aesthetics* Vol. II

- Thomasson, A. (2004) "The Ontology of Art", in Kivy, P. (ed.) (2004).
- Thomasson, A. (2005) "The Ontology of Art and Knowledge in Aesthetics", *The Journal of Aesthetics and Art Criticism*, Vol. 63, N° 3, pp. 221-229.
- Tilghman, B. (2006) *Reflections on Aesthetic Judgement and Other Essays*, Ashgate Publishing Company.
- Trivedi, S. "Music and Imagination", in Kania & Gracyk (eds.) (2011)
- Varèse & Chou Wen-Chung (1966) "The Liberation of Sound", *Perspectives of New Music*, Vol. 5, N° 1, pp. 11-19.
- Wallin, Merker & Brown (eds.) (2001) *The Origins of Music*, MIT Press.
- Walton, K. (1970) "Categories of Art", *Philosophical Review*, Vol. 79, N° 3, pp. 334-367.
- Walton, K. (1988) "The Presentation and Portrayal of Sound Patterns", in J. Dancy, J. M. E. Moravcsik & C. C. W. Taylor (eds.), . Stanford University Press, pp. 237-257
- Weitz (1956), "The Role of Theory in Aesthetics", *The Journal of Aesthetics and Art Criticism*, Vol. 15, N° 1, pp. 27-35.
- Wittgenstein, L. (1986) *Philosophical Investigations* (trad. G. E. M. Anscombe), Oxford, Basil Blackwell.
- Wollheim, R. (1980) *Art and Its Objects*, Cambridge University Press.
- Wolterstorff, N. (2003) "Why Philosophy of Art Can't Handle Kissing, Touching and Crying", *The Journal of Aesthetics and Art Criticism*, Vol. 61, N° 1, pp. 17-27.
- Young, J. O. (1999) "The Cognitive Value of Music", *The Journal of Aesthetics and Art Criticism*, Vol. 57, N° 1, pp. 41-54.
- Young, J. O. (2001) *Art and Knowledge*, London & New York, Routledge.
- Young, J. O. (2012) "Resemblance, Convention and Musical Expressiveness", *The Monist*, Vol. 95, N° 4, pp. 587-605.
- Young, J. O. (2014) *Critique of Pure Music*, Oxford University Press.
- Zangwill, N. (1986) "Aesthetics and Art", *British Journal of Aesthetics*, Vol. 26, N° 3, pp. 257-269.
- Zangwill, N. (1994) "Doughnuts and Dickie", *Ratio (New Series)* VII, 1 June, Basil Blackwell Ltd.

- Zangwill, N. (2001) *The Metaphysics of Beauty*, Cornell University Press.
- Zangwill, N. (2007) *Aesthetic Creation*, Oxford University Press.
- Zangwill, N. (2009) (Replies to Six Critics) *Sztuka I Filozofia*, Issue 35.
- Zangwill, N. (2011) "Music, Essential Metaphor and Private Language", *American Philosophical Quarterly*, Vol. 48, N° 1.
- Zangwill, N. (2012) "Listening to Music Together", *British Journal of Aesthetics*, Vol. 52, N° 4, pp. 379-389.
- Zangwill, N. (2013) "Clouds of Illusion in the Aesthetics of Nature", *The Philosophical Quarterly*, Vol. 63, N° 252, pp. 576-596.
- Zangwill, N. (2014) "Metaphor as Appropriation", *Philosophy and Literature*, Vol. 38, n° 1, pp. 142-152.
- Zangwill, N. (2014) "Music, Metaphor and Aesthetic Concepts", *The Journal of Aesthetics and Art Criticism*, Vol. 72, N° 1, pp. 1-11.
- Zangwill, N. (2015) *Music and Aesthetic Reality*, Routledge.
- Zappa& Occhiogrosso (1990) *The Real Frank Zappa Book*, Pan Macmillan.
- Zemach, E. (1970) "Four Ontologies", *The Journal of Philosophy*, Vol. 67, N° 8, pp. 231-247.
- Zemach, E. (1982) "A Plea for a New Nominalism", *Canadian Journal of Philosophy*, Vol. 12, N° 3, pp. 527-537.
- Zemach, E. (1990) "Nesting: The Ontology of Interpretation", *The Monist*, Vol. 73, N°2, pp. 296-311.
- Zemach, E. (1986) "No Identification Without Evaluation", *British Journal of Aesthetics*, Vol. 26, N° 3, pp. 239-251.
- Zemach, E. (1997) *Real Beauty*, Penn State University Press.
- Ziff, P. (1953) "The Task of Defining a Work of Art", *The Philosophical Review*, Vol. 62, N° 1, pp. 58-78.
- Ziff, P. (1973) "The Cow on the Roof", *The Journal of Philosophy*, Vol. 70, N° 19, pp. 713-723.
- Zuckerkandl, V. (1956) *Sound and Symbol: Music and the External World*, trans. Willard R. Trask, Princeton University Press.

