

THE METAPHYSICS OF INEFFABILITY

Silvia L. Y. N. Jonas

BPhil Thesis in Philosophy
delivered to the University of Oxford in Trinity Term 2010

Supervision: Professor A.W. Moore

The Metaphysics of Ineffability¹

Abstract

The existence of the phenomenon of ineffability enjoys wide recognition, and numerous examples can be found in the contexts of the philosophy of mind, the philosophy of religion, and philosophical aesthetics. At the same time, however, its philosophical significance is largely denied, which explains the lack of literature on the topic. My thesis aims to fill this gap by offering a systematic account of the metaphysics of ineffability. The structure of the thesis is as follows. After providing some core definitions, the paper proceeds to examine what, if anything, could qualify as an ineffable entity. Three candidates are examined ('ineffable truths', 'ineffable content' and 'ineffable knowledge'), the first two of which are rejected, the third of which is endorsed. In the discussion of ineffable truths, I challenge arguments for what I call 'fugitive' and 'excess' propositions. In the discussion of ineffable content, I challenge the notion of nonconceptual content in the contexts of perception and aesthetics. In the discussion of ineffable knowledge, I clarify to what extent arguments for the distinctness of knowledge-how from knowledge-that can help elucidate ineffability. Finally, I endorse A.W. Moore's argument for the existence of ineffable knowledge and apply it to the initial examples of ineffability. I conclude that the phenomenon of ineffability must be explained in terms of ineffable knowledge.

This thesis is dedicated to my parents. I owe every success to their faith in me.

I am also deeply indebted to my grandparents' and family's unconditional support.

I should like to express my special gratitude to the Konrad-Adenauer-Foundation, which provided the funding for my studies at Oxford.

¹ Whoever feels compelled to ridicule attempts to discuss the ineffable should consider this thesis a footnote to a blank piece of paper.

Table of Contents

1	Introduction: Making Sense of Ineffability.....	4
2	Instances of Ineffability	5
3	Definitions	7
3.1	Experience, Content, Representations, Truth-bearers	7
3.2	Ineffability, Indescribability, Expression.....	9
3.3	Expressibility: Possible Extensions of ‘Can’.....	12
4	Ineffability Types	15
5	Analysis: What Exactly is Ineffable?	17
5.1	Ineffable Truths?.....	18
5.1.1	‘The Argument from ‘Fugitive Propositions’	18
5.1.1.1	Linguistic Objections	21
5.1.1.2	Metaphysical Objections	23
5.1.2	Lewisian Excess Propositions	27
5.1.2.1	Objections to the Setup of the Argument	30
5.1.2.2	An Alternative Objection to Premise (4).....	32
5.2	Ineffable Content?	33
5.2.1	Nonconceptual Content in Perceptual Experience	34
5.2.1.1	An Argument from ‘Contradictory Content’	36
5.2.1.2	An Argument from ‘Analog Content’	37
5.2.1.3	‘Fineness-of-Grain’ Arguments.....	38
5.2.2	Nonconceptual Content in Aesthetic Experience.....	41
5.2.2.1	Aesthetic Experience – ‘Trivially Ineffable?’.....	44
5.2.2.2	A Twofold Structure of Awareness?	46
5.2.2.3	Non-inferential Content?	47
5.3	Ineffable Knowledge?	49
5.3.1	Knowledge-how	49
5.3.2	Moorean Ineffable Knowledge	54
5.3.2.1	Knowing Through Being Shown: Moore’s account	54
5.3.2.2	The Objection from Externalism.....	61
5.3.2.3	Applications of Moorean Ineffable Knowledge	66
6	Conclusion: Ineffability is Making Sense	67
7	Bibliography.....	69

1 Introduction: Making Sense of Ineffability

Everyone experiences moments of speechlessness. Sometimes the ineffability we experience in those moments feels important. This raises a question: why do we attach importance to what we are unable to put into words? Surely because in many cases we feel that, despite the impossibility of expressing the ineffable, we gain some sort of insight. Arthur Danto captures the feeling of importance in the following way:

“*What* cannot be told of, we want to know, knowing also that we cannot say: the ineffable is that about which all that is to be said is that nothing more is sayable. Still, one persists in wanting to know, there must be some explanation of that which at once stimulates and frustrates our descriptive impulses: at least we want to know if the difficulties are due to the ineffable itself, or to speech.”²

Someone might say: ‘I don’t associate ineffability with *importance* at all; the state of being speechless is nothing more than a temporary failure of one’s cognitive apparatus!’ However, I think that such statements are inaccurate in at least some cases. One can experience speechlessness of two kinds: speechlessness of a mundane kind and speechlessness of an enigmatic kind. For example, everyone knows the feeling of being unable to express one’s emotions in moments of extraordinary joy or sadness: “Words cannot express how I feel at this moment!” is a popular sentence filling the enunciative gap. However, most people don’t spend a long time afterwards reflecting on their inability to put into words what they felt at those moments. There are other moments of speechlessness, however, which stick with us: we remember those occasions again and again, and continue to feel puzzled by our inability to express what we experienced. For example, we feel that a piece of classical music “tells” us something, or that we understand something in a moment of prayer, but we cannot say what it was that we were “told” or what it was that we understood. A hint of enigma attaches to such moments. This is because, I submit, we feel that there is *something* of importance lying behind the linguistic barrier.

In the following, I want to examine to what extent this intuition is justified. That is, I want to examine if ineffability has to do with how the world is, i.e. whether there are metaphysically substantial entities of some sort which resist expression, or whether one should doubt that there actually *is* anything in the world one intends to express but fails to. This examination requires:

² Danto 1973: 46

- (1) providing a number of *exemplary cases* in which the term ‘ineffability’ can be encountered,
- (2) offering an initial *definition* of ‘ineffability’, provided with regard to a standard of expressibility and in contrast to ‘indescribability’,
- (3) identifying the different *types* of ineffability and thus separating the philosophically interesting from the philosophically uninteresting instances of ineffability,
- (4) *analysing* the philosophically interesting cases with respect to what exactly it is about them that we want to call ineffable (possible candidates being ‘truths’, ‘content’, and ‘knowledge’),
- (5) testing whether all philosophically relevant cases of ineffability can be *reduced* to one concept,
- (6) providing a conclusion, based on the results of (5), as to whether our intuition that, if only we were able to express the ineffable, we would be able to reveal and share an important insight “about the world”, is justified.

It is not an easy task to explore a phenomenon that is, by definition, inexpressible. As Donald Davidson framed it, “so often in philosophy, it is hard to improve intelligibility while retaining the excitement.”³ The goal of this enquiry is to show that the notion of ineffability can indeed be made intelligible while retaining the sense of excitement associated with it.

2 Instances of Ineffability

If we want to understand what ineffability is (or could be), we need to ask two closely related questions. The first one is: *what* can be ineffable? What are the things, or entities, or states, to which the predicate ‘ineffable’ can apply?⁴ The second question is: *how*, i.e. in which ways, are these things, or entities, or states, ineffable? It is not possible to answer one question without answering the other. This is because finding out what can be ineffable requires us to look at actual cases of something being ineffable, and these cases feature different ways in which something can be ineffable. Here are eight different cases in which one might speak of ineffability:

³ Davidson 1973: 5

⁴ With the exception of A.W. Moore’s work, the literature about ineffability tends not to be clear enough about what exactly the term ‘ineffable’ is taken to apply to. The notions of ‘ineffable truths’, ‘ineffable content’, ‘ineffable knowledge’, and ‘ineffable experience’ often seem conflated (cf. for example Raffman 1993; De Clercq 2000).

1. (Trivial Ineffability) It is impossible to express a stone.
2. (Nescient Ineffability) I cannot express what this person's name is (because I don't know it) but you can (because you know it).
3. (Impediment Ineffability) I cannot express what this person's name is (because I'm gagged) but you can (because you are not gagged).
4. (Sensational Ineffability) It is impossible to express my sensation of tasting saffron in such a way that another person will come to know how saffron tastes.
5. (Aesthetic Ineffability) It is impossible to paraphrase the content of a painting, a melody, a poem into literal⁵ language.
6. (Mystical Ineffability) I cannot express my knowledge of G-d.⁶ I know that He exists, but cannot explain what I mean by this.
7. (Contradiction Ineffability) It is impossible to express certain states of affairs without producing a contradiction. Imagine the compulsive Liar uttering the sentence "This sentence is false." A way of capturing the paradoxical state of affairs involved in this scenario is to say that the Liar has uttered a sentence that is both true and false, which is contradictory.
8. (Moorean⁷ Ineffability) It is impossible to express certain pieces of knowledge without saying something false. For instance, many people would say that the sentence "There would be no world if there was nobody to perceive it" is an accurate expression of a deep-founded human intuition, although we have reasons to believe that the sentence is false.

These are eight exemplary cases for which the term 'ineffability' might be considered appropriate. However, as I will argue below, only the latter four of these cases of ineffability are of philosophical interest. Before proceeding to that argument, it is crucial to explicate my terminology.

⁵ It is worth noting here that I take 'ineffability' to be equivalent to 'inexpressible in literal language'. If the content of a melody were expressible in metaphorical language only, it would still count as ineffable according to my definition (see below). The reason for this restriction is that there is a high degree of agreement about the meaning and reference of literal expressions (if there is disagreement – unless it is among philosophers of language – a standard dictionary will resolve it), whereas there can be significant disagreement about the meaning and reference of non-literal expressions. By restricting ineffability to ineffability in literal language, I intend to keep possible disagreement about whether or not some linguistic item *x* qualifies as an expression of some content *y* to a minimum.

⁶ Due to religious reasons I will not spell out His name in full.

⁷ I call this kind of ineffability 'Moorean Ineffability' because it is the kind of ineffability A.W. Moore engages with in his book *Points of View*. His theory involves a particular reading of Wittgenstein's *Tractatus*, including a transcendental interpretation of the saying/showing distinction.

3 Definitions

We cannot discuss the metaphysics of ineffability without a proper definitional framework. The core terms that will occur frequently are: ‘experience’, ‘content’, ‘representations’, ‘truth-bearers’, ‘expression’, ‘expressibility’, and ‘ineffability’. Here are the respective definitions.

3.1 Experience, Content, Representations, Truth-bearers

Experiences are events or streams of events in the consciousness of sentient beings. They are phenomenal in character, which is equivalent to saying that they have a phenomenology. The phenomenology of an experience e is *what it is like* for a sentient being to have experience e . Experiences have either one or two components. Experiences like listening to an argument between two people have a sensational component (we sense the sounds, movements, colours of the speakers) and a content component (we grasp the content of what Jimmy yells at Johnny, etc.). Experiences like eating ice-cream have a sensational component (tasting vanilla, for example) but lack a content component.⁸

‘Content’ is a very general term. We can distinguish between the content of an item like a bucket, a book, or an artwork, and the content of a proposition, the content of a (perceptual) experience, the content of a mental state, etc. In all cases, the term refers to that which either constitutes or is contained in the respective item, proposition, experience, or mental state. For my discussion of ineffability, those kinds of content which can become part of a subject’s mental content will be most relevant. On my view, if a subject S receives a piece of content c through some experience e , c will become part of that S ’s mental content m . m consists of c plus some kind of cognitive classification (e.g. propositional attitude formation, emotional reaction, embedding in a network of beliefs, memories, etc.).

There are two general classes of views concerning the truth-aptness of the contents of experiences: one class of views considers *all of* the contents of experience to have accuracy conditions and therefore, to be truth-evaluable; the other class of views considers *none or only some* of the contents of experience to be truth-evaluable. The former notion – that experiences have accuracy conditions – is often spelled out by arguing that experiences somehow *derive* from the contents of belief: some hold that

⁸ This is not uncontroversial. Unfortunately, the limited space of this thesis does not allow me to provide a comprehensive discussion of this issue here. However, my argument against phenomenal contents in section 5.1.1.2 provides an indication of how I would support the view that sensations lack content.

experiences are acquisitions of belief, some hold that experiences are dispositions to form beliefs (David Armstrong⁹, Daniel Dennett¹⁰); and others hold that experiences are grounds of dispositions to form beliefs. There are also different ways in which the other notion – that experiences don't have accuracy conditions – can be spelled out: experiences can be considered 'raw feels' (Thomas Reid¹¹) or 'sense data' (Frank Jackson¹²), or one can embrace some form of adverbialism (Roderick Chisholm¹³) or naïve realism (Charles Travis¹⁴) about experiences. Unfortunately, I cannot add a discussion of these different views here; suffice it to say that I endorse the view that only some parts of our experiences are assessable for accuracy, namely the content components.¹⁵ The content components of experience are representations. The purely sensational components, on the other hand, I take to be non-assessable for accuracy and non-representational. Thus, I consider some parts of a subject's mental content to be representational, i.e. constituted by representations (propositions, images, etc.), and some parts to be non-representational. Making these distinctions will be crucial for the classification of different cases of ineffability and, more generally, for my discussion of ineffability types (see section 4).

Representations are entities representing the world; they are representations of the input we receive perceptually and cognitively and are, therefore, truth-apt. An entity that is truth-apt is truth-evaluable, i.e. its form allows that it be assigned a truth-value (either 'true' or 'false'). A representation is truth-apt because it has content in virtue of which it is either true or false. Roughly, the content of a true representation is how the world must be if the representation is to be true of the world.¹⁶ If a representation is false, then the world is not how it must be for that representation to be true. There are linguistic and non-linguistic representations. Linguistic representations are propositional. Whenever we form a propositional attitude by holding a belief, thinking a thought, uttering a sentence, we produce a representation of the world. Only propositional attitudes aimed at truth, such as beliefs, convictions, assertions, etc. are representations; other propositional attitudes, such as desires, fears, etc. which are not aimed at truth are not representations of the world. The content of non-linguistic representations such as

⁹ Armstrong 1968

¹⁰ Dennett 1991

¹¹ Reid 2003 [1764]

¹² Jackson 1977

¹³ Chisholm 1957

¹⁴ Travis 2004

¹⁵ Again, my argument in section 5.1.1.2 indicates why I endorse this view.

¹⁶ I adopt this characterization of content from Moore 1997: 280.

paintings or sculptures may be expressed by a propositional representation with the same content. However, the sensational component of non-linguistic representations cannot be expressed propositionally.

Finally, it needs to be clarified what qualifies as a truth-bearer. For some philosophers, ‘true propositions’, ‘truths’, and ‘facts’ are three different terms for one and the same thing. However, I will *not* use these terms interchangeably. I will take facts to be part of the “furniture of the world”; they are those things which make truth-bearers true or false. My argument is based on the minimal assumption that there must be some kind of necessary relation between true propositions on the one hand, and facts on the other.¹⁷ Propositions are truth-bearers of a specific form, i.e. of a form that allows linguistic expression. I refrain from claiming that propositions are the only kind of truth-bearers there are; there may be truth-bearers which are not propositional in form.¹⁸ I will take a representation to be any entity (propositional or non-propositional in form) which, in virtue of its content, is either true or false. A truth, on this picture, is any true representation.

3.2 Ineffability, Indescribability, Expression

I will use the terms ‘ineffability’ and ‘inexpressibility’ interchangeably. Roughly, something is ineffable if it cannot be expressed in language. By contrast, something that is not ineffable is effable and can be expressed in language.

It is important to distinguish ‘ineffability’ from ‘indescribability’. Something that is ineffable may very well be describable. The difference between the two concepts can be illustrated by replacing ‘to express’ with ‘to describe’ in some of the above examples: a stone is clearly describable, although it is ineffable; so is my sensation of tasting saffron, the content of works of art, and my knowledge of G-d. I take the crucial difference to be the following: for someone to grasp an *expression* of an item x means to understand what it would be for x to be true, whereas by grasping a *description* of x someone does not necessarily grasp what it would be for x to be true.

Besides distinguishing ineffability from indescribability it is necessary to provide an initial definition of ineffability. The natural way of defining ineffability is in contrast to its opposite, expressibility. Yet, in order to understand what the *modal* term ‘expressibility’

¹⁷ Even though I am sympathetic to truth-theories that involve at least some form of correspondence between truth-bearers and reality, my argument is neutral with regard to which theory of truth should be adopted.

¹⁸ For an argument against the view that all truth-bearers are propositions see Thomson 1969.

means, we first need to define what constitutes an *actual* instance of expression. A.W.

Moore suggests the following definition of ‘ x expresses y ’:

“ x expresses y if and only if (i) x is a linguistic item with content that makes it either true or false, (ii) y is a non-linguistic item with content that makes it either true or false, and (iii) the content of x entails the content of y .”¹⁹

‘ x expresses y ’ thus pertains to the state of affairs in which, for a given non-linguistic item y , there is a corresponding linguistic item x whose content entails the content of y . Note that, according to this definition, we can speak of a case of expression if the content of a linguistic item x *entails* the content of a non-linguistic item y . The content of x entails the content of y , but it is not necessarily *equivalent* to the content of y ; there is an asymmetrical relation between x and y . This simply means that not only the sentence ‘Lemons are yellow’ but also the sentence ‘Lemons are yellow and oranges are orange’ qualifies as an expression of the proposition that lemons are yellow.²⁰ Note also that ‘ x expresses y ’ is defined as a property of truth-apt entities, such as representations. It is not defined as a property of entities which are not truth-apt, such as stones (more on this in section 4). However, Moore does not elaborate on how he understands the notion of a ‘linguistic item with content’. I think that this is problematic because, as Moore’s definition stands, all kinds of odd scenarios qualify as cases of expression. A dialogue in Lewis Carroll’s *Through the Looking Glass* illustrates my point:

“‘When I use a word,’ Humpty Dumpty said in rather a scornful tone, ‘it means just what I choose it to mean—neither more nor less.’
 ‘The question is,’ said Alice, ‘whether you CAN make words mean so many different things.’
 ‘The question is,’ said Humpty Dumpty, ‘which is to be master—that’s all.’”²¹

The question is: what is a linguistic item? Should a word that was invented by Humpty Dumpty and that can only be understood by Humpty Dumpty qualify as an expression? Or should it be at least possible that other people come to understand it? An answer to the question what qualifies as a genuine expression depends on how we understand ‘language’, and, in a derivative sense, who or what we take to be capable of producing genuine expressions. Moore indicates a range of possible answers:

“[...] can a truth be expressed linguistically only if there is a sentence in some existing language that has precisely that truth as its content? Or does it suffice that there be some (true) sentence in an existing language that has that truth as *part* of its content? If the latter, in what sense of ‘part’? Does it

¹⁹ Moore 2003a: 173

²⁰ The relevance of this qualification will become clear in section 5.1.2.

²¹ Carroll, Lewis 1994 [1872]: 98ff

perhaps suffice that there be some sentence in an *extension* of an existing language that has that truth as its content? If so, then what counts as an extension of an existing language? Or is the appeal to existing languages too restrictive? Does it suffice that there be some (true) sentence in some *possible* language that has that truth as (part of) its content? Then what counts as a possible language?”²²

The standards I will adopt are quite low.²³ On my view, a genuine expression is a piece of a possible language (it doesn’t have to be a piece of an actual language). However, a genuine expression must, in principle, be capable of being understood by more than one finite being. I am consciously restricting attention to finite beings, so when I talk about expression, I mean what can be expressed, in the sense of the above definition, by a finite being. I deliberately rule out infinite beings and that’s purely stipulative. Why the stipulation? First of all, if expressions had to be capable of being understood by infinite beings only, it would expand the notion of expressibility in such a way that it wouldn’t cover the intuitive cases I think it should cover (namely those which concern the kind of ineffability human beings experience). Moreover, there just is a very interesting distinction to be drawn between what can be expressively achieved using finite resources and what would be possible for an infinite being. In this context, it seems to me that this is a much more interesting distinction to be drawn than any of the distinctions that might be drawn between all sorts of different finite beings.²⁴

When I say that an expression must be ‘in principle capable of being understood’, I mean that an expression must be communicable, i.e. shareable through some process of learning. This allows for Humpty-Dumpty-style word inventions so long as the meaning of such an invented expression be learnable by other finite beings. The qualification rules out private languages because they cannot, by definition, be learned by someone else.

Someone might worry that different groups of individuals might possess conceptual schemes so different from one another that they resist intertranslation, such that a linguistic item which qualifies as an expression of *y* in one (linguistic) community has no counterpart in another (linguistic) community, and as a result, *y* would be effable in one community but ineffable in another. However, this worry can be alleviated with an

²² Moore 2003b: 162 (his emphases)

²³ I should note that I don’t think my standards are the ‘right’ ones, and all other standards would be ‘wrong’; there is no question of right or wrong when it comes to definitions. However, I do hope that my definitions cover at least some (but hopefully many) of our intuitions.

²⁴ As will become clear in the course of my thesis, the distinction between finite and infinite beings will turn out to be irrelevant as well. This is because I will argue that ineffability is a matter of ineffable knowledge, and ineffable knowledge could neither be expressed by finite nor by infinite beings (see section 5.3.2.1). Nevertheless, I think it is important to discuss all these possible distinctions at this point, in order to address as many of the intuitions the reader might have about expressibility as possible.

argument by Davidson. He refutes the doctrine according to which there are or could be conceptual schemes so different from one another that they resist intertranslation. Conceptual schemes are described as

“ways of organizing experience; [...] systems of categories that give form to the data of sensation; [...] points of view from which individuals, cultures, or periods survey the passing scene.”²⁵

If, as conceptual relativists argue, there are conceptual schemes so different from one another that an intelligible concept in one scheme cannot be translated into an intelligible concept in another scheme, then it seems that

“the beliefs, desires, hopes and bits of knowledge that characterize one person have no true counterparts for the subscriber to another scheme.”²⁶

Davidson argues that neither the notion of a complete failure of translatability nor of a partial failure of translatability of conceptual schemes is intelligible. To see why, he asks the reader to consider how we decide whether a conceptual scheme is an adequate one. A conceptual scheme qualifies as adequate if it ‘fits the facts’ or ‘fits the totality of possible sensory evidence’. These two expressions are nothing else but a circumlocution for ‘being true’. Davidson notes:

“Our attempt to characterize languages or conceptual schemes in terms of the notion of fitting some entity has come down, then, to the simple thought that something is an acceptable conceptual scheme or theory if it is true. [...] And the criterion of a conceptual scheme different from our own now becomes: largely true but not translatable. The question whether this is a useful criterion is just the question how well we understand the notion of truth, as applied to language, independent of the notion of translation. The answer is, I think, that we do not understand it independently at all.”²⁷

Applying Davidson’s argument to our concerns, we can conclude that, if there is a linguistic item which qualifies as a suitable expression of a given content *y* in one conceptual scheme, it is in principle possible to transport that expression into another conceptual scheme by translation.

3.3 Expressibility: Possible Extensions of ‘Can’

Having defined an *actual* instance of expression by defining ‘*x* expresses *y*’, we can now extend our definition to *possible* instances of expression: ‘*y* is expressible if it *can* be expressed by some *x*’. However, there is a problem with the word ‘can’ which makes the content of this definition less clear than it at first appears. The problem is that we tacitly

²⁵ Davidson 1973: 5

²⁶ Davidson 1973: 5

²⁷ Davidson 1973: 16

give a context-defining interpretation to the word ‘can’. In an article on the foundations of mathematics, Friedrich Waismann touches on the problem of reading the word ‘can’ adequately. He raises the problem of how to establish that two sets *can* be paired-up in a one-to-one relation.²⁸ A possible answer would be that a set of cups and a set of spoons are equal in number if we can place exactly one spoon into each cup. Waismann then asks: what if the cups and the spoons are in different drawers and it is therefore impossible to place the spoons in the cups?, and provides the answer himself:

“One will reply that this was not the intention of the explanation; it does not depend on whether I actually place the spoons in the cups but whether I *can* place them in the cups. Very well! But what does the expression ‘I can’ mean here? Is it that I have to be physically able to distribute the spoons among the cups? This would be entirely uninteresting. Obviously, what we wish to say is that I can distribute the spoons among the cups because there are just as many samples of both sorts. That is, in order to recognize whether the correspondence is possible, I must already know that the sets are numerically equivalent. [...] The statement: ‘The two sets *can* be associated to one another,’ is being reinterpreted into the statement which is entirely distinct from it: ‘The two sets *are* associated to one another,’ which means, ‘There is actually a relation which permits such a correspondence.’”²⁹

Waismann makes a crucial point in this section: the word ‘can’ (more precisely: the words ‘can be’) is tacitly reinterpreted as ‘are’; a modal claim is read as an actual fact. I think that Waismann’s worry applies to all occurrences of ‘can’ where it is not clear from the context how ‘can’ should be interpreted.³⁰ When we say that ‘*y* is expressible iff it can be expressed by some *x*,’ it is not clear from the context how ‘can’ should be interpreted; so it requires clarification.

There are several different kinds of possibility which the word ‘can’ can express: logical, metaphysical, nomological possibility come to mind.³¹ Something, say *P*, is logically possible if and only if no contradiction can be derived from *P* (perhaps in conjunction with certain definitions) using the standard rules of deductive inference. Something is nomologically possible for a relevant body of (e.g. physical or biological) laws, if and only if *P* is consistent with the body of truths entailed by those laws.

²⁸ Waismann discusses Frege’s definition of the concept ‘number’ by means of a definition of equinumerosity. Frege argues that two sets are equal in number if they can be paired-up in as one-to-one relation. Waismann points out that this attempt to establish that two sets are numerically equivalent is circular.

²⁹ Waismann 1951: 109 (his emphasis)

³⁰ If I say: “Tiffany is too weak to lift the stone but Samson can lift it”, it is clear from the context that ‘can’ refers to Samson’s physical ability and not to some other circumstances (e.g. that Samson can lift the stone because he happens to be at Tiffany’s today).

³¹ There are far more than three kinds of possibility (e.g. epistemic possibility, conceptual possibility, *de dicto* possibility, *de re* possibility, human possibility, technical possibility, etc.). However, the three kinds I picked will suffice to make my point. For elaborate discussions on this topic see Gendler/Hawthorne 2002.

Metaphysical possibility is usually taken as a primitive, although it can be described in terms of ‘how things might have been’.³² These different kinds of possibility are usually taken to differ in strength. Metaphysical possibility is a stronger notion of possibility than logical possibility because metaphysical possibility *implies* logical possibility. The reverse does not hold, i.e. logical possibility does not imply metaphysical possibility. Likewise, nomological possibility implies metaphysical possibility (and thus also logical possibility), but metaphysical possibility does not imply nomological possibility. Why is this relevant to our discussion of ineffability?

It is relevant because our definition of ‘expressibility’ depends on the kind of possibility we take the word ‘can’ to signify. Thus, our definition of expressibility ‘*y* is expressible iff it can be expressed by some *x*’ can be read as (a) ‘*y* is expressible iff it is logically possible that there be some *x* which expresses *y*’; or it can be read as (b) ‘*y* is expressible iff it is metaphysically possible that there be some *x* which expresses *y*’; or it can be read as (c) ‘*y* is expressible iff it is nomologically possible that there be some *x* which expresses *y*’. (a), (b), and (c) denote three very distinct extensions of ‘can’. If we restrict our definition of expressibility to the nomological extension of ‘can’, a non-linguistic item *y* is expressible if it is nomologically possible for there to be a linguistic item *x* which expresses *y*. However, what if we are dealing with a highly complex non-linguistic item, call it *HC*, whose expression is so complicated that it would take more than a human lifetime to express it? Then it would certainly be nomologically impossible to express *HC* and it would consequently count as ineffable. However, declaring a non-linguistic item ineffable simply because of the complex structure its expression would require seems counterintuitive. Intuitively, there is also a clear sense in which *HC* is expressible, namely if only humans happened to live longer. Restricting our definition of expressibility to the nomological extension of ‘can’ seems too restrictive and is thus not a satisfactory definition.

What if we go for (c), the logical extension of ‘can’? On that definition, *HC* would surely not count as ineffable because it is certainly logically possible that there be a being with a longer life-span than humans which manages to express *HC*. However, I think that (c) is so permissive that it becomes vacuous; I cannot see how anything would turn out ineffable on that definition. It is always logically possible that there be an expression of any non-linguistic item, unless, of course, the non-linguistic item is already *defined* as ineffable – in this case, it is logically impossible to express it, but then it is also not an

³² Gendler/Hawthorne 2002: 4f

interesting kind of ineffability but a merely definitional one. It is not a case which provides any insight into the nature of ineffability.

In order to capture all the interesting instances of ineffability, our definition of expressibility should neither be vacuous nor too restrictive. It therefore seems to me that (b) employs the right extension of ‘can’: ‘ y is expressible iff it is metaphysically possible that there be some x which expresses y ’. On that definition, a non-linguistic item like *HC* would not count as ineffable merely because of the complexity of its expression.

Having clarified the difference between ineffability and indescribability, provided a definition of ‘ x expresses y ’ and explained the way in which the word ‘can’ is to be interpreted the sentence ‘ y is expressible if it can be expressed by some ‘ x ’, we can now define ineffability as follows:

- ❖ Definition (*Ineffability*): A non-linguistic item y is ineffable if and only if it is metaphysically impossible that there be a linguistic item x whose content a user of any finite language would recognize as entailing the content of y and whose content is, in principle, communicable to other finite beings.

With this definition in place, we can now proceed to the discussion of the different types of ineffability, by means of which we will separate the philosophically interesting from the philosophically uninteresting cases of ineffability.

4 Ineffability Types

In section 2, I described eight different examples of ineffability: ‘trivial’, ‘nescient’, ‘impediment’, ‘sensational’, ‘aesthetic’, ‘mystical’, ‘contradiction’ and ‘Moorean’ ineffability. I also stated that only the last four of these cases are interesting. The results of the discussion of the possible extensions of ‘can’ will now enable me to explain why.

According to my definition, a non-linguistic item is ineffable if and only if it is metaphysically impossible that there be a linguistic item x whose content a user of a finite language would recognize as entailing the content of y . This definition immediately rules out ‘nescient’ and ‘impediment’ cases. Even if I cannot express what your name is because I don’t know it or because I am gagged, it clearly remains metaphysically possible for it to be expressed (e.g. by some person which knows what your name and is who isn’t gagged). So what your name is is not ineffable according to the definition I provided.

What about the case I called ‘trivial’ ineffability, the impossibility to express a stone? A stone is a non-linguistic item, but is it metaphysically impossible that there be a linguistic item whose content a user of any finite language would recognize as entailing

the content of the stone? Admittedly, this is a rather puzzling question: what is the ‘content’ of a stone? *Does* a stone have content at all? I have defined the ‘content’ of an item as that which either constitutes or is contained in the respective item. The content of a stone thus seems to be the stone itself. So is it metaphysically possible that there be a linguistic item whose content a user of any finite language would recognize as entailing the stone? I believe not. An item like a stone is not the kind of content that could be entailed by an expression; expressions don’t entail physical items. Rather, expressions entail contents which are apt to become part of a subject’s mental content. So a stone is ineffable according to my definition. However, is that an interesting kind of ineffability? It doesn’t seem like it. The reason is, I think, that we would only find something interestingly ineffable if we *expected* it to be capable of becoming part of a subject’s mental content but then find out that it actually isn’t. We never expected a stone to be expressible, so we are not intrigued or puzzled by our inability to express it.³³ Thus, cases of trivial ineffability are genuine cases of ineffability yet philosophically uninteresting.

The fourth case was what I called ‘sensational ineffability’. The example I used was ‘It is impossible to express my sensation of tasting saffron’. It is important to be clear about what exactly this is supposed to mean. It is *not* supposed to mean that I cannot *describe* my sensation of tasting saffron. Of course I can describe that sensation as ‘spicy’, ‘slightly sweet’, ‘a bit like soap’, etc. The point is, however, that such a description would not suffice to make someone who has never tasted saffron before understand my sensation, i.e. make him understand what it tastes like. This holds for sensual experiences in general. Our ability to talk about sensual experiences is based on the fact that we *share* those experiences, i.e. that other people have had the same sensual experiences. Thus, when I say ‘This tastes like saffron’, my friend who knows what saffron tastes like will have an understanding of what I am talking about which is similar to mine. My other friend who has never tasted saffron won’t have an understanding similar to mine, even though he probably understands what ‘saffron’ refers to. Thus, unless we are talking about sensual experiences which our interlocutors have had as well, sensations are ineffable in the sense that we cannot instil the same experiences in someone merely by expressing them. However, is the ineffability of sensations philosophically interesting?

³³ Someone might argue that by expecting a stone to be expressible, we are simply making a category mistake: a stone is not the kind of thing that can be expressed, and that’s all there is to say. However, all things which are ineffable according to my definition are ‘not the kinds of thing that can be expressed’ – yet we wouldn’t be making a category mistake if we expected them to be expressible. As I will argue below, there are states of knowledge which are ineffable, but given that they are states of knowledge, and given that we normally expect states of knowledge to be expressible, there is no sense in which expecting a state of ineffable knowledge to be expressible could constitute a category mistake.

Let's see how sensations fit into the framework given by my definition of ineffability. A taste is a non-linguistic item. Is it metaphysically possible that there be a linguistic item whose content a user of any finite language would recognize as entailing the content of the taste? *Does* a sensation like taste have content at all? I distinguished above between two different experiential components, a phenomenological component and a content component. The experience of tasting saffron certainly has a phenomenology: there is something that it is like to taste saffron. However, it doesn't have content in the sense described above. In that sense, the taste of saffron and the stone have something in common. And just like in the case of the stone, I believe that the ineffability we experience with regard to sensations is not a philosophically interesting kind of ineffability; we don't normally expect others to be able to learn what saffron tastes like by mere description. Language is simply not a kind of medium through which we receive sensual input. Having a sensation like tasting requires being *acquainted* with what is tasted to become part of a taster's mental content, just like having a sensation of what red looks like requires being acquainted with something red. Yet, we are not puzzled by this.³⁴

So far, I have argued the following. Using my definition of ineffability, I have ruled out those possible candidates for entities that can be ineffable which are only in practice, but not in principle, ineffable: 'nescient' and 'impediment' ineffability. I have distinguished entities which are in principle ineffable according to a criterion of expectation and argued that the philosophically interesting cases are those in which an entity that we *would* expect to be effable *turns out* to be ineffable. Thus I have also ruled out 'trivial' and 'sensational' ineffability, and we are left with the four remaining cases: 'aesthetic', 'mystical', 'contradiction' and 'Moorean' ineffability. I will now proceed to examine what, if anything, their ineffability consists in.

5 Analysis: What Exactly is Ineffable?

How are we supposed to make sense of the intuition that, despite the impossibility of expressing the ineffable, we gain some sort of insight through it? One way of understanding this statement is by claiming that we grasp an ineffable truth. The

³⁴ It is also interesting to note that instances of ineffability differ with regards to where the felt ineffability directs our *attention*. Whereas having a sensation like tasting draws our attention "inward", i.e. towards something that is happening inside of my body, the ineffability of aesthetic or religious experiences draws our attention outward, toward "something" we feel is "out there". Although this is not enough to make a meaningful distinction between interesting and uninteresting kinds of ineffability, it may be a good indicator for why we are puzzled by some instances of ineffability and not by others: instances of ineffability that direct our attention inward is less likely to invoke the feeling that we have just "come to know" something important about the world than instances of ineffability that direct our attention outward.

predicate ‘true’ is most commonly applied to propositions³⁵ (and all other entities taken to express propositions, such as utterances, sentences, etc.). However, on the assumption that propositions constitute content, one could also argue that ‘truth’ is a property of contents. In the following two sections I will argue that both the notion of ‘ineffable truths’ and the notion of ‘ineffable contents’ are incoherent and should thus be dismissed.

5.1 Ineffable Truths?

For some, the very question whether there are or could be ineffable truths involves a contradiction in terms. I have in mind those philosophers who hold the opinion that truth is nothing more than a “trivially transparent property” of all propositions which are instances of the equivalence schema “The proposition that *p* is true if and only if *p*”.³⁶ On this picture, propositions are by definition expressible, and truth is by definition a property of propositions only; so the notion of ineffable truths is ruled out by definition. Evidently, this position contains two substantial presuppositions, i.e. that all propositions are expressible and that truth is a property of propositions alone. If one or both of these presuppositions is challenged, the notion of ineffable truths regains its appeal.

In this section, I will present and reject two kinds of argument for ineffable truths. The first one I will refer to as the argument from ‘fugitive propositions’, the second one I will refer to as the argument from ‘Lewisian excess propositions’.

5.1.1 The Argument from ‘Fugitive Propositions’

There are propositions, call them *fugitive propositions*, whose expression is impossible even though we may have some idea of what the content of the proposition must be, and what its truth entails. As John Hawthorne and Graham Oppy put it, these are “propositions that no language user – and perhaps no possible language user – is capable of grasping.”³⁷

What could such fugitive propositions be? I think one can identify at least three different candidates for fugitive proposition, differing from each other with respect to the cause of their fugitiveness. The first kind of fugitive proposition could be called *discovery-fugitive*. Here is an example. Given that science (physics, for instance)

³⁵ Even though most philosophers agree that truth is a property of propositions, there is a controversy about the nature of propositions: some suggest that they are sets of possible worlds, others suggest that propositions are structured entities with constituents (although there is no agreement as to what sort of things these constituents are, and as to what binds the constituents of propositions together).

³⁶ Cf. Horwich 1998

³⁷ Hawthorne/Oppy 1997: 186

continuously discovers truths about the world, it is reasonable to assume that there are truths “waiting” to be discovered in the future. For example, one of the chief goals of the Large Hadron Collider at CERN is to generate evidence for or against the hypothesis that there exists an elementary particle called the ‘Higgs boson’. The research conducted at CERN is based on the assumption that there is a *truth* about the existence of the Higgs boson which can be discovered: either, ‘The Higgs boson exists’ or, ‘The Higgs boson does not exist’. As matters stand now, we don’t know which one of the two sentences expresses the truth about the Higgs boson; however, we *do* know that one of the two sentences *does* express the truth about the Higgs boson. Therefore, even though we may not know which of the two sentences is the true one, *discovery-fugitive* propositions are not ineffable in the sense stipulated by our definition.

The second kind of fugitive proposition could be called *finitude-fugitive*. It is caused by a restriction common to all finite beings: we cannot express infinite propositional conjunctions because expressing them would take an infinite amount of time. It may be perfectly possible to express each single instance of the infinite collection, but it will always be impossible to express its conjunction, i.e. to express all single instances at the same time. Examples of such propositions are the conjunction of all empirical true propositions of the past, present and future³⁸, the conjunction of all propositions expressing instances of the law of the excluded middle, and, more generally, the conjunction of all propositions representing instances of a schema. However, as Bertrand Russell noted in ‘The Limits of Empiricism’, the impossibility of writing down the complete decimal expansion of π (and, more generally, of performing infinitely many tasks in a finite time) is not a logical impossibility but a merely ‘medical’ one.³⁹ In other words, this candidate for fugitive propositions arguably represents a particular case of ‘impediment ineffability’. The impossibility of expressing infinite conjunctions as a finite being can arguably be considered a kind of physical constraint. It is true that this kind of constraint could not be removed as easily as in the above example where the person was gagged; however, the example still seems not very interesting philosophically because we have a clear grasp of what it is that prevents us from expressing *finitude-fugitive* propositions – namely, our own finitude.

The third kind of fugitive proposition could be called *perspectival-fugitive*. In ‘What is it like to be a bat?’, Thomas Nagel famously states that a successful account of the mind-

³⁸ Hawthorne/Oppy 1997: 186

³⁹ Russell 1935: 143f. The claim that performing infinitely many tasks in a finite time is a merely medical impossibility has given rise to some controversy; cf. Moore 1989.

body-problem must be able to accommodate the subjective character of conscious experience.⁴⁰ In *The View from Nowhere*, he states that

“[...] not everything there is can be gathered into a uniform conception of the universe from nowhere within it. If certain perspectives evidently exist which cannot be analyzed in physical terms, we must modify our idea of objective reality to include them. If that is not enough, we must admit to reality some things that cannot be objectively understood.”⁴¹

In his view, reductionist accounts of the mind, like physicalism, fail to acknowledge the fact that providing an objective account requires both an abstraction from and an accommodation of something that could be called ‘perspectival facts’. His argument can be rendered as follows:⁴²

1. Just as human beings have conscious experience, other organisms, for example bats, have conscious experience as well.
2. For a bat to have conscious experience means that there is *something* it is like to be a bat. (Nagel refers to this ‘something’ as “the subjective character of experience” (p. 436); “the phenomenological features of experience” (p. 437); “[f]acts about what it is like to be an *X*” (p. 437); “facts that embody a particular point of view” (p. 441).
3. The human sensory apparatus is unsuitable to form representations of such non-human perspectival facts. Consequently, humans can neither describe nor form conceptions of such perspectival facts – they are “beyond the reach of human concepts” (p. 441).

Nagel speaks of *facts* beyond the reach of human concepts. Given that facts *obtain*, i.e. that they are *true* of the world, it seems that Nagel is arguing for a particular kind of ineffable truths. He writes:

“Reflection on what it is like to be a bat seems to lead us, therefore, to the conclusion that there are facts that do not consist in the truth of propositions expressible in a human language. We can be compelled to recognize the existence of such facts without being able to state or comprehend them.”⁴³

⁴⁰ Nagel 1974

⁴¹ Nagel 1986: 163

⁴² Page references are to his (1974).

⁴³ Nagel 1974: 441

A similar line of thought, which has come to be known as the ‘Knowledge Argument’, is provided by Frank Jackson. He argues for the existence of ‘qualia’, which are taken to be the phenomenal aspects of our mental lives, as follows:

“I think that there are certain features of the bodily sensations especially, but also of certain perceptual experiences, which no amount of purely physical information includes. Tell me everything physical there is to tell about what is going on in a living brain, the kind of states, their functional role, their relation to what goes on at other times and in other brains, and so on and so forth, and be I as clever as can be in fitting it all together, you won't have told me about the hurtfulness of pains, the itchiness of itches, pangs of jealousy, or about the characteristic experience of tasting a lemon, smelling a rose, hearing a loud noise or seeing the sky.”⁴⁴

In order to illustrate his point, he asks the reader to imagine Mary, a poor creature confined to a black-and-white room, who (through black-and-white books and a black-and-white TV) learns everything physics can tell us about the world.⁴⁵ Specifically, she learns all physical facts there are to learn about colour perception, including facts about wavelengths, retina receptors, etc. A physicalist would have to say that she knows all there is to know about the world, and specifically about colour perception. However, there is an intuitive sense in which Mary seems to learn something *new*, something *in addition* to the physical facts, once she is finally freed from her black-and-white prison and sees the colour red for the first time. Jackson argues that the intuition that Mary acquires an extra piece of knowledge when seeing red for the first time is strong enough to reject physicalism. I take Jackson’s argument to be intended to establish a similar, if not identical, point as Nagel’s argument: there is a perspectival-fugitive proposition about what it is like for Mary to see the colour red.

These *perspectival-fugitive* propositions are the kind of ineffable truths I will focus on. In the following, I will raise two principal objections to the notion of ineffable truths, one linguistic, one metaphysical, based on which I will argue that the notion of ineffable truths is incoherent.

5.1.1.1 Linguistic Objections

A linguistic objection against the notion of ineffable truths is that we cannot make sense of the concept of a ‘truth’ independently of the concept of ‘linguistic expression’,

⁴⁴ Jackson 1982: 127

⁴⁵ Jackson 1986: 291

or better: that any truth is expressible in one way or another, and therefore, that any truth is expressible linguistically. The argument, which is due to A.W. Moore, runs as follows.⁴⁶

- 1) For any truth t , there is at least one possible corresponding true representation r of t . For example, if I say that there is a truth in the witness' testimony, this either means that the witness' testimony contains a true sentence (or utterance), or that there is the possibility of someone making a claim which is entailed in one way or another by the testimony, that is true. There cannot be a truth in the witness' testimony unless some such possibility of expression obtains. That possibility *just is* the possibility of the truth's being expressed.
- 2) A possible corresponding true representation r of t is identical to a possible (not necessarily linguistic) expression e of t .
- 3) Thus, for any truth t , there is a possible (not necessarily linguistic) expression e of t .
- 4) Representations have content which answers to reality and in virtue of which they are either true or false.
- 5) For any (arbitrarily chosen) true representation r_1 and for any (arbitrarily chosen) true representation r_2 , there must be an integrating possible true representation r_3 whose content is either identical to, or at the least includes the contents of r_1 and r_2 taken together. This holds on the assumption that representations are made true by a single, unified reality.
- 6) Since there are radical differences of type⁴⁷ and content between some representations, an integrating representation like r_3 can only integrate r_1 and r_2 abstractly. Imagine a true representation produced by a bumble-bee when seeing a flower, and a true representation produced by a biologist when observing a clownfish defend his territory.⁴⁸ On the assumption that both representations are true of a single, unified reality, there will be some content "between" these two representations. This content will supply a truth which, like any other truth, is expressible. However, given the differences in kind and content between r_1 and r_2 , the integrating representation r_3 must be sufficiently abstract. That is, in order to encompass the contents of two true representations r_1 and r_2 , an integrating

⁴⁶ Moore 2003b: 163

⁴⁷ The type of a representation is determined by the role it plays in the psychology and phenomenology of the subject that produces it if it informs a belief of that very subject (cf. Moore 1997: 283).

⁴⁸ We could have chosen any two arbitrary truths.

true representation r_3 must abstract from the points of view from which r_1 and r_2 were produced.

- 7) A sufficiently high degree of abstraction can be provided only by linguistic expressions.
- 8) Therefore, for any true representation r corresponding to truth t , there is a linguistic expression e_L which expresses r .

The crucial assumption here is that every truth can possibly be represented, and the key move of the argument is to show that every representation can be expressed linguistically.

5.1.1.2 Metaphysical Objections

Besides linguistic considerations, there are also metaphysical worries about the notion of ineffable truths. A.W. Moore, for example, devotes an entire chapter of *Points of View* to the refutation of what he calls the ‘Specious View’.⁴⁹ The core intuition of this view is that perspective, in addition to being a feature of representations of reality, is a feature of reality *itself*. The *Specious View* can be characterized as follows:

“*The Specious View*: Perspective is a characteristic not only of representations but also of what is represented. There are perspectival features of reality, which figure in perspectival facts. Perspectival facts are like any other facts in that for them to obtain is for the world to be a certain way. But they are unlike other facts in that their obtaining is itself relative to a point of view.⁵⁰ The perspectival facts that obtain from one point of view are different from those that obtain from another. What makes (some) true perspectival representations true is, precisely, the obtaining of perspectival facts.”⁵¹

Arguably, as we have seen above, one adherent to this view is Nagel. According to his argument for the existence of perspectival facts, these facts (e.g. facts about how the world looks/sounds/feels from specific perspectives) obtain relative to a point of view. This distinguishes perspectival facts from “ordinary” facts (e.g. facts about how many siblings Lisa has) because ordinary facts do not obtain relative to specific points of view only; ordinary facts obtain no matter from which point of view they are considered. According to the adherent to the ‘Specious View’, true perspectival representations are made true by the obtaining of perspectival facts.

⁴⁹ Moore 1997: 41ff

⁵⁰ A point of view is “a location in the broadest possible sense”; they “include points in space, points in time, frames of reference, historical and cultural contexts, different roles in personal relationships, points of involvement of other kinds, and the sensory apparatuses of different species.” (Moore 1997: 6)

⁵¹ Moore 1997: 42

A.W. Moore refutes the ‘Specious View’ by arguing that it leads to an incoherent conception of reality. I will employ the following example to sketch Moore’s refutation. Imagine two people, Ernie and Bert, standing in the vicinity of Carfax Tower in the centre of Oxford. Ernie is standing 200 metres away of the tower in front of Christchurch College. Bert is standing one metre away, right at the bottom of Carfax Tower on Cornmarket Street. Ernie says: ‘Carfax Tower is 200 metres to the right.’ Bert says: ‘Carfax Tower is one metre to the left.’ The adherent to the *Specious View* would be committed to say that both representations are true because of the obtaining of their respective perspectival facts. However, if what Ernie says is true because it is a fact that Carfax Tower is 200 metres to the right, and if what Bert says is true because it is a fact that Carfax Tower is one metre to the left, then Carfax tower is both 200 metres to the right and one metre to the left, which is contradictory.

A related line of argument can be applied to Nagel’s argument for perspectival facts about what it is like to be an *X*. Instead of assuming a location-dependent perspectival fact as in the example above, we assume a perspectival fact dependent on a specific sensory apparatus. Imagine a bat using its sonar technique, thus producing a representation of a tree through sound propagation. Let’s call the bat’s representation of the tree r_B . Imagine further a human being looking at the tree and, through her human sensory system, producing a representation of the tree, call it r_H . Let’s suppose that the content of both representations is that there is a tree. The adherent to the *Specious View* would have to say that r_B is true because it corresponds to a perspectival fact obtaining for the bat, call it f_B , and that r_H is true because it corresponds to a perspectival fact obtaining for the human, call it f_H . Consequently, although both representations have the same content, they are made true by different facts – a paradoxical result.

The way adherents to the *Specious View* usually try to remove this contradiction is by claiming that perspectival facts are only “accessible” from the respective points of view of bats, humans, etc. However, it is difficult to make sense of what exactly this restricted accessibility comes down to. One possibility would be to claim that the conjunction of all perspectival facts obtaining from a specific point of view constitutes a world. In that case, however, the adherent to the *Specious View* would be committed to the existence of as many worlds as there are points of view – a number far exceeding the number of possible worlds, since every possible world “contains” infinitely many points of view. However, even if the adherent of the *Specious View* were ready to accept such a vast number of worlds, there remains another problem.

The other problem is that he would have to tell some story about the relations between the “small” perspectival worlds to the “big” world (i.e. that one world constituted by all those facts which obtain from any point of view, the world of physics and mathematics), since every person is an “inhabitant” both of the “big” world and of at least one (but probably many more) of the “small” perspectival worlds. Simply denying the existence of such a “big world” is not an option, at least not for anyone who believes that true representations are made true by one and the same unified reality.⁵² Let’s go back to Ernie and Bert and their respective claims that Carfax tower is 200 metres to the right and one metre to the left. In order to resolve the blatant contradiction, Ernie and Bert must find a way to integrate both their representations into one picture, and the way to do this is by transcending their subjective points of view. A way of transcending their points of view in order to accommodate their contradictory beliefs about the location of Carfax Tower is to point out that Ernie and Bert both inhabit the same world but occupy different spatial locations of that world. In order to interpret Ernie’s claim that Carfax Tower is 200 metres to the right correctly, we must take into account Ernie’s spatial location. The same goes for Bert’s claim that Carfax Tower is one metre to the left. If we do this, we will see that both of their claims are true relative to their spatial location. The point is that, for any number of true representations, it must be possible, not only for each single one of them, but for all of them taken together, to be true. In other words: there must be a single, unified world of which all true representations are true. Whoever denies this rules out the possibility of meaningful disagreement: if there wasn’t a single, unified world, it wouldn’t even be possible to identify a contradiction in Ernie’s and Bert’s claims.

Another problem for the adherent to the *Specious View* is that he would need to explain the individuation of, and the relations between, the different perspectival worlds. This would prove to be difficult since there is arguably both overlap of the perspectival worlds and interaction between them. For example, although Ernie’s and Bert’s perspectival worlds differ with regard to the truths about Carfax Tower, they are invariant (and thus overlap) with regard to other truths, e.g. about physical laws (if $E=mc^2$ is true in Ernie’s world, it is true in Bert’s world as well). Does this overlap mean that Ernie and Bert “share” a world with regard to physical laws but live in separate

⁵² If someone were ready to deny that representations are made true by a single, unified reality, then my argument would not stand (although I am not aware of anyone who holds such a view). However, a consequence of denying that representations are made true by a single, unified reality would be that meaningful disagreement about the truth or untruth of representations would no longer be possible.

worlds with regard to Carfax Tower? This seems highly absurd. The adherent to the *Specious View* could try to argue that there isn't strictly speaking any overlap of perspectival worlds but rather some kind of *correspondence* between the truths in one world and truths in other worlds. On this view, truth T_E ($E=mc^2$) is a truth of Ernie's world and corresponds to truth T_B ($E=mc^2$) which is true in Bert's world. This view avoids the problem of overlap. However, it also creates new ones. For example, the adherent to the *Specious View* would have to explain what the correspondence relation between truths like T_E and T_B consists in. A natural answer would be that it consists in the fact that T_E and T_B have the same content (namely, $E=mc^2$). However, they don't. T_E and T_B are true of two different worlds and therefore cannot have the same content. I don't see how a correspondence relation could be explained without some sort of appeal to content, and therefore, I consider the strategy of invoking correspondence relations between perspectival worlds unconvincing. Given that many perspectival worlds are incompatible (imagine the perspectival worlds constituted by the sum of all facts obtaining from the perspective of Ernie and Bert), the enterprise of trying to accommodate all perspectival worlds in such a way that interaction between them is possible and mutual exclusion is unproblematic seems bound to fail.

Another attempt for the adherent to the 'Specious View' to save his theory from contradiction involves embracing some form of truth-relativism, i.e. some form of the view that truth is relative to frameworks such as points of view, perspectives, etc. However, as Plato noted in the *Theaetetus*,⁵³ solutions of this kind suffer from the defect of self-refutation: either, the claim that all truths are relative to a specific framework is itself a truth relative to a specific framework (and thus not universally applicable); or, the claim that all truths are relative to a specific framework constitutes an exception to the rule that all truths are relative to a specific framework, in which case one could readily ask how this exception can be motivated and whether there are further exceptions to the rule. The adherent to the *Specious View* could, of course, cheerfully admit that he is operating in a particular framework which may be inaccessible to people not working in that framework. However, as noted above, it seems to me that this isn't really an option if we want to allow for the possibility of meaningful disagreement between the adherent to and the opponent of the *Specious View*.

To sum up, the notion of ineffable truths as inferred from perspectival facts must be rejected for two principal reasons. One reason is that the notion of ineffable truths is

⁵³ Plato 1961 [428-348 BCE]: 170e-171c (p 876f)

linguistically incoherent: the notion of a ‘truth’ implies the possibility of linguistic expression. The other reason is that the attempt to accommodate the perspectival facts corresponding to ineffable truths in one’s metaphysics leads to an incoherent picture of reality. Moreover, the doctrine of perspectival facts entails relativism about truth, which is a self-refuting doctrine.

5.1.2 Lewisian Excess Propositions

Another way of arguing for ineffable truths can be extracted from an argument for excess propositions provided by David Lewis. In *On the Plurality of Worlds*, Lewis argues that propositions are sets of possible worlds. On this view, any subset of possible worlds constitutes a proposition which would be expressed by a thought or sentence which was true exactly in every world of that subset and in no other world. A Lewisian proposition thus carves the set of all possible worlds into two classes, so that a sentence expressing the proposition is true in all and only one half of worlds and false in all and only the other half of worlds.

Before getting into the examination Lewis’ argument, it is important to note the following. The fact that this section deals with a notion of expressibility defined in terms of the *equivalence* of expressions and propositions marks a departure from the definition of expression I gave above. Recall that, according to clause (iii) of my definition, it would suffice for a sentence *S* to count as an expression of proposition *P* if *S* *entailed* *P* (in which case *S* would not be true in *all* the possible worlds which constitute *P* but only in a specific *subset* of those worlds). My definition does not require that *S* be *equivalent* to *P*. According to my definition, the sentence ‘Lemons are yellow and oranges are orange’ counts as an expression of the proposition that lemons are yellow. Lewis’ argument, on the other hand, assumes a stronger notion of expression. For him, a sentence *S* expresses a proposition *P* just in case *S* is *equivalent* to *P*. I will explain in due course why it is important to keep this distinction in mind.

Lewis bases his argument for excess propositions on a paradox, known as ‘Kaplan’s Paradox’, which emerges as follows:

1. We assume that the cardinality of the set of all possible worlds is K .
2. If propositions are sets of possible worlds, then each subset of the set of all possible worlds constitutes a proposition.
3. Since a set of size K has 2^K subsets, there are 2^K propositions. 2^K is strictly greater than K .

4. Every proposition can possibly constitute the sole content of a being's thought at a given time.
5. For each such scenario, there is a possible world.
6. Therefore, there are at least 2^K possible worlds. This contradicts (1).⁵⁴

Since Lewis wants to defend the notion that propositions are sets of possible worlds, he needs to find a way to avoid this paradox. He chooses to avoid it by denying (4), i.e. by holding that not every but only a very specific, very small subset of all propositions can possibly constitute the content of a being's thought. All other propositions, although they exist, *cannot possibly* constitute the content of a being's thought. Therefore, they can also never be expressed and are thus, ineffable propositions.⁵⁵ They are 'excess propositions', as it were.

Lewis thus suggests restricting possible thought contents to a small number of propositions only. He proposes the following functionalist restriction criterion:

"A man or a beast or a god, or anything that is a thinker at all, has a thought with a certain content in virtue of being in a state which occupies a certain functional role. This definitive functional role has to do with the causal relations of that state to the thinker's sensory input, his behavioural output, and his other states."⁵⁶

His idea, then, seems to be this. Imagine any thinking being *B*. *B*'s having a thought is directly dependent on *B*'s being in a state with a certain functional role. In other words, the functional role of *B*'s state determines the content of *B*'s thought. Thus, Lewis claims, "there can be only as many different possible contents of thought as there are different definitive functional roles."⁵⁷ What does Lewis understand by a 'functional role'? He explains them as "the relevantly different ways of thinking".⁵⁸ According to Lewis' theory of content,⁵⁹ an assignment of content to states must *fit* the functional role of that state. Whether a content fits the functional role of a state depends on certain principles of rationality. For example, if a person *A* is in a state *s* which is assigned content *c* (where *c* consists of some system of propositional attitudes such as beliefs, desires, etc.), then the assignment of *c* to *s* fits the functional role of *s* iff *s* tends to produce conduct which is conducive to the aim of the relevant propositional attitude.

⁵⁴ Lewis 1986: 104f

⁵⁵ Someone might argue that we could imagine a machine which produces expressions of those propositions which cannot ever constitute the content of a being's thought. However, on my account, the machine would not be producing genuine expressions since they could not possibly be understood by anyone else. See the discussion of the possible extensions of 'can' above.

⁵⁶ Lewis 1986: 106

⁵⁷ Lewis 1986: 106

⁵⁸ Lewis 1986: 106

⁵⁹ Cf. Lewis 1986: 27ff

Lewis thus argues, in line with standard rational choice theory, that we can assign beliefs and desires to people according to their behaviour. If we assume that people are rational and that rational beings tend to act to satisfy their preferences by maximising the expected utility of their actions, then it is likely that rational beings have a preference ordering amongst possible actions. What we know about human behaviour provides a fairly full picture of this preference ordering. On this picture, it is possible to employ some mathematical representation theorem (like Savage's representation theorem⁶⁰) to conclude that an action is rational if and only if, relative to the subject's beliefs and desires, it has the highest subjective expected utility of all available options. From this rational preference structure of actions, we can now assign relevantly belief-like things (credences such that the higher the credence, the more believed the proposition/the higher the degree of belief) and relevantly desire-like things (utilities or values assigns to outcomes, such that the higher the utility, the more desired the outcome) to rational agents. On this way of characterizing propositional attitudes, allowing an arbitrary proposition to be the sole content of a belief will require a respective arbitrary stream of behaviour. However, assuming that we don't want to accept arbitrary streams of behaviour as rational, we must rule out all propositions belief in which couldn't be assigned a stream of rational behaviour. Hence, we must rule out all propositions which consist in arbitrary sets of possible worlds.

Put less abstractly: if I am in a certain state to which a certain content consisting of the desire expressed by 'I want ice-cream' is assigned, then that content fits my state iff my state tends to produce action conducive to my getting ice-cream (for example, sending my father to the nearest ice cream parlour). The content would *not* fit my state if it tended to produce action conducive to something other than getting ice-cream (for example, going to the gym). However, the tricky thing is that there are plenty of absurd content assignment scenarios which qualify as 'fitting' according to this criterion. For example, rather than thinking that I can get ice-cream at the next door ice-cream parlour, we could construe a completely absurd assignment, for example that I desire a cactus and, due to a strange yet valid chain of inductive reasoning supported by all my previous experiences with ice-cream parlours, I expect the ice-cream parlour to serve me cactus. In order to rule out such absurd yet valid assignments, Lewis suggests a 'principle of humanity' for content assignments, i.e. a principle designed to somehow rule out

⁶⁰ Cf. Savage 1972

unreasonable cases. It is those ruled out content assignments which are, according to Lewis, examples for inexpressible propositions. He elaborates:

“They are ineligible contents of thought because they are utterly unpatterned and miscellaneous. I suggest that an unthinkable content is one that can never be correctly assigned because, whenever it fits the functional roles of the thinker’s states, some more favoured content also fits.”⁶¹

At this stage, we need to go back to the two different notions of expression I pointed out at the beginning of this section. The notion of ineligible contents of thought, which I converted into the notion of ineffable propositions, depends on Lewis’ definition of expression, according to which an expression must be *equivalent* to a proposition. On my definition of expression, the following holds. An expression need only *entail* a proposition. The sentence ‘Lemons are yellow and oranges are orange’ expresses the proposition that lemons are yellow and oranges are orange because it picks out a set of possible worlds in which lemons are yellow and oranges are orange. However, given that the set of all possible world in which lemons are yellow and oranges are orange is a subset of the set of all possible worlds in which lemons are yellow, the expression ‘Lemons are yellow and oranges are orange’ also entails the set of all possible worlds in which lemons are yellow, and therefore, counts as an expression of the proposition that lemons are yellow. It is now possible to argue along the same lines that a Lewisian ineffable propositions, call it LIP, does not come out ineffable on my definition. Arguably, it is possible that there be an expression picking out a subset of the set of all possible worlds constituting LIP which entailed the set of all possible worlds constituting LIP. So arguably, Lewisian excess propositions don’t come out ineffable on my definition which suggests that there is much less of an issue in Lewis’ argument than might at first appear. Nevertheless, I think it is important to consider Lewis’s argument as an argument for ineffable propositions because there may be some philosophers who reject clause (iii) of my definition of expression, but also because it is interesting to see what happens if expression is defined in terms of equivalence rather than entailment: then it looks as if Lewis is presenting us with a challenge. So let’s see how Lewis’ argument might be attacked.

5.1.2.1 Objections to the Setup of the Argument

There are several reservations one might have about the setup of Lewis’ argument. After all, hardly anyone fully endorses Lewis’ metaphysical framework. The first

⁶¹ Lewis 1986: 108

objection which could be brought forward is against the notion that propositions are *sets* of possible worlds, rather than, say, disjunctions of ways things might be. However, that objection won't help to refute Lewis' argument. The reason is that Kaplan's paradox does not only arise if propositions are taken to be sets of worlds; it arises for anyone who believes that, for any disjunction of ways worlds could be, there is a proposition which is true in all and only those worlds.

One could also argue that propositions slice logical space more thinly than possible worlds by considering distinct propositions picking out the exact same set of possible worlds. Consider for example the propositions 'Louis' kite is rectangular' and 'There is a 90-degree-angle between all four sides of Louis' kite'. Arguably, both propositions are true in exactly the same possible worlds. However, given that they provide slightly different pieces of information, they are distinct. However, even if such arguments falsify Lewis' premise that for K possible worlds, there are 2^K propositions, it doesn't get rid of Kaplan's paradox. Actually, it only makes things worse by multiplying the number of propositions and, thereby, the number of worlds.

Another possibility to object to the setup of Lewis' argument is to claim that the cardinality of the set of possible worlds is not K but 1 because there is only one actual world. However, that wouldn't help either. The paradox is not a paradox about *worlds* but a paradox about *descriptions of possibilities*. Hence, it also arises if the cardinality of worlds is assumed to be 1 and possibilities are taken to be maximal descriptions ('ersatz worlds') of ways things might be. If the assumption is that the cardinality of the set of maximal descriptions of the way things might be is D , then there are 2^D non-maximal descriptions of ways things might be, and thus, 2^D possible settings in which a lonely thinker thinks one of those non-maximal descriptions. It seems that, to avoid the paradox, one either has to deny that there are such maximal descriptions of possible worlds, or deny that they can be thought, as Lewis does. I agree with Lewis that denying the existence of maximal descriptions of the way things might be is not the right way to go. This is because, as I mentioned above, meaningful disagreement is only possible if we assume the unity of both the actual world and each possible world (or maximal description of possibilities), and thereby, the consistency of the respective associated propositions. Thus, the way to avoid Kaplan's paradox is to find a way to deny that completely arbitrary propositions can constitute the content of a thinker's thought.

A rather simple solution is to deny that unthinkable propositions are propositions at all. If a proposition is defined as the kind of thing that could be thought or uttered, then

arguably those excess propositions Lewis invokes won't even count as propositions because they are neither thinkable nor utterable. This would entail that all propositions *can* be thought. However, given that Lewisian excess propositions are nothing but (arbitrary) disjunctions of descriptions of ways things might be, it seems difficult to make sense of the idea that a 'non-arbitrary' disjunction of propositions would qualify as a proposition while an 'arbitrary' disjunction wouldn't. Again, this does not seem to be the right way to object to Lewis' argument.

5.1.2.2 An Alternative Objection to Premise (4)

Lewis attacks premise (4) of Kaplan's paradox, i.e. the claim that every proposition can possibly constitute the sole content of a being's thought at a given time. His attack can be converted into an argument for ineffable propositions. However, maybe he could have challenged premise (4) in a way that did not depend on the strange notion of ineffable propositions? Maybe one could challenge premise (4) by arguing that it is impossible that a being thinks exactly one proposition at a given time.

Recall that Kaplan's paradox arises because it is assumed that every proposition can possibly constitute the *sole* content of a lonely thinker's thought at a given time. However, on second thought, it seems rather dubious that a thinker can think one proposition alone (i.e. have one propositional attitude) at a given time. Take beliefs for example.⁶² Arguably, all beliefs we hold are embedded in a network of other beliefs. We hold beliefs on the basis of other beliefs, we know that belief b_1 entails belief b_2 , we make inferences from one belief to another, etc. Therefore, one could argue that holding one belief always implies that one holds several other beliefs at the same time (in the "back" of one's mind, as it were) and thus, that lonely-thinker-scenarios aren't real possibilities. Arguments of this kind supposedly show that there are not as many worlds as Kaplan's paradox seems to establish. However, I think this strategy can easily be corrupted. It may well be that we never actually think about one proposition only at a given time. However, we can simply reformulate Kaplan's paradox in terms of a lonely writer, i.e. someone who writes down exactly one proposition at a given time. This would re-establish that, for K possible worlds, there are 2^K propositions the lonely writer could write down, and thus, 2^K possible worlds we can infer from those scenarios.

I have considered several different ways of attacking Lewis' argument, but it seems that it cannot be resisted if the underlying definition of expression is one of equivalence

⁶² Something similar can be argued with respect to other propositional attitudes such as desires.

rather than entailment. This is not only an interesting philosophical result, but also a motivation for adopting a definition of expressibility in terms of entailment rather than equivalence.

In the preceding sections, I have argued that ineffability cannot be explained in terms of ineffable truths because the concept ‘ineffable truths’ is incoherent by considering Nagel’s idea of fugitive propositions and providing counter-arguments against it. I have also discussed Lewisian excess propositions, which only turn out ineffable on Lewis’ definition of expressibility (which is stronger than mine), and argued that we can get around them by assuming a definition of expression which requires expressions to entail, rather than be equivalent to, the propositions they express. Lewis’ argument for ineffable propositions can thus be avoided by adjusting one’s definition of expression. By rejecting that ineffability is a matter of ineffable truths or ineffable propositions, I have provided reasons to believe that we don’t experience ineffability because of some sort of ineffable entities (facts, propositions) which are part of the “furniture of the world”. I will now proceed to examine whether ineffability can be explained in terms of ‘ineffable content’.

5.2 Ineffable Content?

In order to examine the notion of ‘ineffable content’, it will be useful to look at the philosophical discussion of nonconceptual content. The notion of nonconceptual content is primarily deployed in the context of three different explanatory projects: the characterization of the contents of perceptual experience, the characterization of the contents of aesthetic experiences, and the characterization of the contents of subpersonal (subdoxastic) states. I will focus on the first two.

The debate is relevant to my discussion of ineffability because it challenges the widely held view that a person’s mental contents are functions of the concepts she possesses. This view puts a *conceptual constraint* on mental contents. If it were true that all of a person’s mental contents are functions of the concepts she possesses, then, on the assumption that we can express linguistically everything that we can conceptualize, all of a person’s mental contents would arguably be expressible in words. If, on the other hand, it could be shown that a person’s mental contents are at least partly nonconceptual, then this could be employed as support for the hypothesis that there are ineffable mental contents.⁶³

⁶³ Many proponents of nonconceptual content consider nonconceptual content to be representational. I have argued above that everything representational must ultimately be expressible as well. The reason why the debate about nonconceptual content is still relevant for my purposes is that it focuses on the question

The central idea behind the notion of nonconceptual content is that it is possible for someone to be in a mental state with content, yet without possessing the concepts needed to express the content of that state. Following Christopher Peacocke, I take concepts to be “constituents of those intentional contents which can be the complete, truth-evaluable, contents of judgment and belief.”⁶⁴ Besides ‘ordinary’ concepts like ‘red’ or ‘wooden’, I take it that there are also demonstrative concepts like ‘this’ or ‘that’, and indexical concepts like ‘here’ or ‘there’. Concepts are individuated according to their content, so that two concepts C and D are distinct

“if and only if there is some completing content Σ such that the complete content $\Sigma(C)$ is distinct from the complete content $\Sigma(D)$; or, in other words, if and only if there is some content $\Sigma(C)$ such that someone for whom the question arises can rationally judge $\Sigma(C)$ without judging $\Sigma(D)$.”⁶⁵

Finally, I take it to be uncontroversial that any content that can be expressed by means of a proposition is conceptual.

The case of nonconceptual content constitutes a challenge to theories claiming that the way in which the world is conveyed to a sentient being depends on that being’s conceptual capacities. Proponents of nonconceptual content hold that conceptual and nonconceptual content are two different types of content, distinguished by the fact that only one type, *viz.* conceptual content, can be brought into propositional form whereas the other one, nonconceptual content, cannot. I will now proceed to examine whether ineffability can be explained in terms of nonconceptual content.

5.2.1 Nonconceptual Content in Perceptual Experience

José Luis Bermúdez states that, “[a]lthough few proponents of the idea of nonconceptual content explicitly adhere to an ineffability claim, the debate about the relation between perceptual content and demonstrative concepts is best seen as a debate about ineffability.”⁶⁶ The debate revolves around the conceptual constraint which conceptualists impose on all mental contents and which nonconceptualists intend to lift. Accounts of nonconceptual perceptual content differ with regards to how they attempt to lift the conceptual constraint. The conceptual constraint can be lifted ‘globally’, so that a subject’s mental contents come out as entirely independent of the concepts she

whether *the subject* in the state with nonconceptual content can express the content of its state. Thus, even if nonconceptual content were representational and thus expressible for, say, a scientific observer, it could still be useful to explain instances of ineffability.

⁶⁴ Peacocke 2001a: 243

⁶⁵ Peacocke 2001a: 243

⁶⁶ Bermúdez 2007: 66

possesses, or ‘locally’, so that a subject’s mental contents come out as partly independent of the concepts she possesses. Possible worlds semantics, according to which propositions are functions from worlds to truth values, constitute a global lifting of the conceptual constraint. However, on that view, all mental contents *trivially* come out as nonconceptual because even propositional contents are identified with sets of worlds. For this reason, I will set global lifting strategies aside and focus on ‘local’ lifting strategies. According to those, only some parts of a person’s mental content are not a function of the concepts she possesses. I will now examine several accounts according to which the content of perceptual states is not exhausted by its propositional content.⁶⁷ On those views, mental content is partly composed of ineffable content.

Note that there is a difference between mere sensations and nonconceptual perceptual content:

“What distinguishes the theory of nonconceptual content from theories of the sensational component of perceptual experience is that nonconceptual content is representational. If a given perceptual experience has a nonconceptual content that Φ , this means that the experience represents Φ as holding in the world. Conceptual and nonconceptual contents are distinguished not by whether they are representational but according to how they represent. They are distinguished according to whether, in specifying how they represent the environment as being, we need to restrict ourselves to concepts possessed by the perceiver. This is in stark contrast to the traditional distinction between sensation and belief, according to which the sensational component has no role to play in explaining how the experience represents the world as being.”⁶⁸

So there is a sense in which nonconceptual content is ‘representational’. But then, why is it interesting for our argument? Didn’t I argue above that all representations are ultimately linguistically expressible? I have two answers. First: yes, I argued that all representations are expressible, but I didn’t say that they had to be expressible by the producer of the representation. A representation could be so difficult to express linguistically that a human being would need, say, a very sophisticated neurological machine to form the expression. In that case, the producer of the representation would be *in practice* (though not *in principle*) incapable of expressing the representation and thus experience it as ineffable. Second: it isn’t at all clear that my use of ‘representational’ and

⁶⁷ See Bermúdez 2007 for a brief overview of the different positions of the debate. I am ignoring arguments which respond to an ambiguity pointed out by Richard Heck (2000), namely that the conceptual/ nonconceptual distinction can either be seen as a distinction between two different types of content (‘content-view’) or between two different types of content-bearing state (‘state-view’). It seems to me that there is wide agreement on the unmotivatedness and untenability of the state view (cf. Bermúdez 2007: 68)

⁶⁸ Bermúdez 1998: 50

Bermúdez use of ‘representational’ coincide. After all, he speaks of two different kinds of representation which are correlated to the two different kinds of content. We will need to find out what exactly is meant by that.

5.2.1.1 An Argument from ‘Contradictory Content’

Tim Crane infers the existence of nonconceptual content from the alleged fact that human experience can represent contradictory content. He uses a phenomenon known as the ‘motion after-effect’ as an example:

If you stare for a period of time at a scene which contains movement in one direction, and then turn your attention to an object in a scene which contains no movement, this object will appear to move in the opposite direction to that of the original movement. [...] But the above description is not quite right. For although the stationary object *does* appear to move, it does not appear to move relative to the background of the scene. That is, there is a clear sense in which it also *appears to stay still* [...].⁶⁹

Crane’s description is supposed to demonstrate that the content of some of our mental representations can be contradictory, since it is logically impossible for something to be both moving and not moving. On the assumption that conceptual content must be consistent – and I think this is a sensible assumption since inconsistent conceptual content would be incapable of realization in a possible world – Crane concludes that contradictory content must be nonconceptual.

However, Crane’s argument is flawed. The reason is that, in order to determine that the content of one’s experience is contradictory because it represents an object as both moving and not moving, one needs to employ the concept ‘moving’. As Hugh Mellor points out, if we didn’t possess the relevant concepts, we wouldn’t be able to detect the contradiction.⁷⁰

Now Crane could argue that that is the whole point, i.e. the minute we bring concepts into play, we are playing the wrong game. As long as content is taken to be conceptual, it has to be consistent. But since experience tells us otherwise, we should assume that content is nonconceptual. However, I still don’t think that this is a way to establish the existence of nonconceptual content. We simply don’t experience inconsistent content. We receive content through perception, and sometimes (as in the Waterfall case), when we try to describe the content we receive, we end up using contradictory concepts. However, that doesn’t establish that the content itself is

⁶⁹ Crane 1988: 142 (his emphasis)

⁷⁰ Mellor 1988: 147ff

contradictory. It thus seems to me that arguments for the existence of nonconceptual content invoking the contradictoriness of content are bound to fail. Yet there are other ways of arguing for nonconceptual content.

5.2.1.2 An Argument from ‘Analog Content’

Fred Dretske has argued for a difference between ‘analog content’ and ‘digital content’, the former characterizing the contents of perception, the latter characterizing the contents of propositional attitudes.⁷¹ He illustrates his view with the following example:

“I see a red apple in a white bowl surrounded by a variety of other objects. I recognize it as an apple. I come to believe that it is an apple. The belief has a content that we express with the words, ‘That is an apple.’ The content of this belief does not represent the apple as red, as large, or as lying next to an orange. I may have (other) beliefs about these matters, but the belief in question abstracts from the concreteness of the sensory representation (icon, sensory information store, experience) in order to represent it simply as an apple. However, these additional pieces of information are contained in the sensory experience of the apple.”⁷²

Dretske argues that the transition from the sensory representation of an apple (by seeing it) to the cognitive representation of an apple (by forming a belief about it) is characterized by “a systematic stripping away of components of information [...] in order to feature *one* component of the information”.⁷³ He calls this procedure of exposing one particular component to the exclusion of all others ‘digitalization’ and argues that all propositional attitudes, being conceptual, are necessarily digital representations. The perceptual experience itself, on the other hand, constitutes an ‘analog’ representation which carries much richer information than any proposition could express. Thus, it is concluded, analog content must be nonconceptual.

However, it seems to me that Dretske’s description of propositional attitude formation on the basis of rich (analog) perceptual content is not convincing as an argument for the existence of nonconceptual content. The reason is that it seems at least possible for every piece of information (or component of analog content) to become the object of our selective attention and thus, to feature in a propositional attitude with digital content. Moreover, as Armstrong points out, Dretske’s account seems compatible with the idea that perception is *nothing else* but the constant acquisition of (propositional)

⁷¹ Dretske 1981: Chapter 6

⁷² Dretske 1983: 61

⁷³ Dretske 1983: 61 (his emphasis)

beliefs about one's environment, even if those beliefs fade away almost immediately.⁷⁴ In Dretske's apple-scenario, for example, it is arguably possible that we acquire many more beliefs than 'That is an apple', including beliefs about the specific features of that very apple ('The apple is half-red half-green'; 'The apple has a dent right next to its stalk'), about the close environment of the apple ('The apple is in a white bowl'; 'The white bowl is made of porcelain'), etc. It may be the case that we never pay conscious attention to these beliefs because we focus on one specific belief only. Yet it may very well be that perception consists in forming numerous beliefs, all of which could *potentially* figure in a propositional attitude. Thus, simply postulating surplus perceptual content will not suffice to make a case for nonconceptual content.

5.2.1.3 'Fineness-of-Grain' Arguments

The account of nonconceptual content which has provoked most discussion in recent years is known as the 'fineness of grain' argument and has initially been triggered by Gareth Evans. In *The Varieties of Reference*, he raises the following question: "Do we really understand the proposal that we have as many colour concepts as there are shades of colour that we can sensibly discriminate?"⁷⁵ Evans' question is a rhetorical one, suggesting that we can perceptually discriminate many more shades of colour than we can conceptualize. Peacocke picks up on Evans' question and argues that our conceptual resources are insufficient to capture the 'fineness of grain' of the perceptual input we receive (or in other words, that our discriminative abilities outstrip our conceptual abilities).⁷⁶

Opponents of nonconceptual content, like John McDowell, search for conclusive ways of aligning our conceptual abilities with our discriminative abilities. For example, an opponent of nonconceptual content could argue for the existence of 'demonstrative concepts' which are exactly as fine-grained as that which we receive perceptually. In this vein, McDowell argues that conceptual capacities determine perceptual experiences. He suggests a strategy which Michael Luntley calls the 'kidnapping strategy'. The key move of this strategy is to invoke something that could be called 'short-lived recognitional capacities' in order to capture the exact content of a person's perceptual experience. Thus, McDowell writes:

⁷⁴ Armstrong 1983: 65

⁷⁵ Evans 1982: 229

⁷⁶ Peacocke 1992: 67ff (Peacocke argues for two kinds of non-conceptual content, 'scenario content' and 'proto-propositional content'. Both are arguments for the claim that the fineness of grain of our perceptual experiences outstrips our conceptual resources.)

“It is possible to acquire the concept of a shade of a colour, and most of us have done so. Why not say that one is thereby equipped to embrace shades of colour within one’s conceptual thinking with the very same determinateness with which they are presented in one’s visual experience, so that one’s concepts can capture colours no less sharply than one’s experience presents them? In the throes of an experience of the kind that putatively transcends one’s conceptual powers – an experience that *ex hypothesi* affords a suitable example – one can give linguistic expression to a concept that is exactly as fine-grained as the experience by uttering a phrase like ‘that shade’, in which the demonstrative exploits the presence of the sample.”⁷⁷

McDowell thus argues that recognizing a certain content implies having the concept to recognize it.⁷⁸ By invoking a higher-order sortal like ‘shade’, and combining it with a demonstrative like ‘that’, we can give expression to whichever shade we want by using the concept ‘that shade’. At this point one could sceptically ask whether we can count such demonstrative capacities as concepts at all. Yet McDowell has an answer for this:

“In the presence of the original sample, ‘that shade’ can give expression to a concept of a shade; what ensures that it is a concept – that thoughts that exploit it have the necessary distance from what would determine them to be true – is that the associated capacity can persist into the future, if only for a short time, and that, having persisted, it can be used also in thoughts about what is by then the past, if only the recent past. What is in play here is a recognitional capacity, possibly quite short-lived, that sets in with the experience.”⁷⁹

At the same time, however, one need not be in possession of a linguistic expression for a concept in order to possess it. An important aspect highlighted by this argument is that for something to be expressible, the subject does not need to use a canonical expression; as long as the subject has some kind of demonstrative resource (“this”, “that”) to express a piece of non-conceptual content, it counts as expression.

However, Peacocke points out that the contender of nonconceptual content should not use a supplemented demonstrative like ‘that shade’ as an example to make his case. The demonstrative ‘that’ should not be supplemented by a general concept like ‘shade’ if it is to capture the exact fineness-level of the relevant experience. This is because supplemented demonstrative concepts already cut too finely:

“Both the person who has the concept scarlet, and has an experience whose content he expresses using the phrase ‘that scarlet’, and someone who merely

⁷⁷ McDowell 1994: 56f

⁷⁸ Note that nothing in this argument runs counter to the claim that our capacities for discrimination are sometimes enhanced through the acquisition of new concepts (cf. Bermúdez 2007: 62) on discrimination enhancement through training).

⁷⁹ McDowell 1994: 57. McDowell’s strategy is obviously not confined to colours but can be employed for all kinds of fine-grained perceptual content.

has the more general concept red, can be seeing a given scarf as having the same finely-individuated shade.”⁸⁰

Thus, the conceptualist should rather use unsupplemented demonstrative concepts like ‘this’ or ‘that’ to make sure every nuance of the perceptual experience is covered by the demonstrative. Yet despite this minor correction of McDowell’s account, Peacocke acknowledges that invoking such demonstratives rules out fineness-of-grain arguments for perceptual content:

“Since these unsupplemented perceptual-demonstratives exist, and can pick out fine-grained properties, the anti-conceptualist should not try to rest his case on fineness of grain.”⁸¹

Far from giving up his conviction of the existence of nonconceptual perceptual content, Peacocke then submits that an argument for nonconceptual content must rest on an argument from animal perception. He argues as follows:

“If the lower animals do not have states with conceptual content, but some of their perceptual states have contents in common with human perceptions, it follows that some perceptual representational content is nonconceptual.”⁸²

Imagine a human being, Alf, in a perceptual state with content that Alf would express through the sentence ‘There is a mushroom’ and a cow in a perceptual state with the same content. Arguably, Alf and the cow are in states with the same content. Alf has the concepts to express it, the cow presumably doesn’t. If we adopt Peacocke’s line of argument, we are forced to conclude that Alf is in a perceptual state with nonconceptual content (because the content of his state can be shared by the cow). However, given that he takes it for granted that animals never have states with conceptual content, it seems to me that the success of this argument depends on assuming very specific conditions for concept-possession. Concept possession must be taken to be a matter of being potentially able to *express* one’s mental content. Cows and humans clearly differ with regard to their expressive capacities. That is a very strong way of defining concept possession, however. Someone could equally well argue that it is sufficient for a person to possess the concept *F* that she be able to *discriminate* *F*s from the rest of the perceptually received content. But if concept possession comes down to nothing more than being able to discriminate *F*s from *non-F*s, then Peacocke’s argument fails. This is because, arguably, both animals and humans are able to discriminate *F*s from *non-F*s

⁸⁰ Peacocke 2001b: 610

⁸¹ Peacocke 2001b: 611; In his (2001a), Peacocke argues that recognitional capacities must be distinguished from demonstrative capacities because the former rely on memory capacities in a way in which demonstrative capacities do not. However, I don’t think this point is relevant in this context.

⁸² Peacocke 2001b: 614

(e.g. mushrooms from trees) and thus, there is a sense in which also animals possess concepts. Unless Peacocke provides a plausible argument to the effect that concept possession requires more than being able to discriminate *F*s from *non-F*s, his argument from animal perception fails to make a good case for the existence of nonconceptual perceptual content.

5.2.2 Nonconceptual Content in Aesthetic Experience

Examining ineffability in the context of aesthetic experience means looking at a very specific kind of content circumscribed to a very particular experiential domain. Someone might be worried that this domain is too narrow to capture all the different instances of ineffability I named at the beginning of this paper. My answer to that worry is that, if ineffability turned out to be explicable in terms of aesthetic content, the account could possibly be extended to different instances of ineffability as well. Thus, it is worth considering ineffability in the context of aesthetics.

Works of art are expressive. They seem to communicate content by invoking an aesthetic experience in the beholder/auditor.⁸³ Most people feel that the content of a specific work of art can only be transported through that very artwork, or in other words, that an artwork cannot be “paraphrased” into language without remainder. If, for example, a friend invited us to enjoy Maria Abramovic’s *The Artist is Present* together, we would find ourselves rather disappointed if we found out that our friend didn’t mean visiting the *Museum of Modern Art* but reading a description of the performance. Likewise, we would find it odd if someone told us that, instead of listening to Chopin’s *Boléro*, we could equally well read a review by an acknowledged critic who claims to be paraphrasing it. I take it to be commonsensical that the content we receive through an aesthetic experience *seems* inextricably linked to the form in which the content is presented, where the form of an artwork could be characterized as “the ensemble of choices intended to realize the point or purpose of an artwork”.⁸⁴ Therefore, it seems impossible to render a painting, a melody, a poem, a metaphor in literal language. In that sense, content transported through aesthetic experience is ineffable. Wittgenstein phrases the point as follows:

“I should like to say: ‘These notes say something glorious, but I do not know what.’ These notes are a powerful gesture, but I cannot put anything side by side with it that will serve as an explanation. A grave nod. James: ‘Our

⁸³ Carroll, Noel 2006: 89

⁸⁴ Carroll, Noel 2006: 78

vocabulary is inadequate.’ Then why don’t we introduce a new one? What would have to be the case for us to be able to?”⁸⁵

Wittgenstein’s question seems justified: if we cannot express something conveyed through an artwork, why don’t we invent the appropriate vocabulary? Defenders of aesthetic content would reply that the nature of aesthetic content is such that, in virtue of its nature, it precludes being captured in words. In other words, that aesthetic content is (at least partially) nonconceptual. The goal in this section is to examine whether we can make sense of such a *specific kind* of aesthetic, nonconceptual content communicated through aesthetic experience. If the notion is coherent, ineffability could perhaps be explained in terms of aesthetic content.

I have already ruled out the possibility of ineffable representational content (see point 5.1.1.1 and 5.2.2). Thus, if there is such a thing as ineffable aesthetic content, it must be non-representational. To clarify some terminological points: I take works of art to be the objects of aesthetic experiences and the beholder/auditor of a work of art to be the subject of an aesthetic experience.

Here are some reasons philosophers have given for assuming that aesthetic content is a specific kind of content which can only be communicated through the respective artworks and which cannot be rendered in language (I will henceforth refer to the proponents of such a view as ‘aestheticians’ and to opponents of this view as ‘anti-aestheticians’).

“If all meanings could be adequately expressed by words, the arts of painting and music would not exist. There are values and meanings that can be expressed only by immediately visible and audible qualities, and to ask what they mean in the sense of something that can be put into words is to deny their distinctive existence.”⁸⁶

“I think every work of art expresses, more or less purely, more or less subtly, not feelings and emotions which the artist *has*, but feelings and emotions which the artist *knows*; his *insight* into the nature of sentience, his picture of vital experience, physical and emotive and fantastic. [...] Such knowledge is not expressible in ordinary discourse. The reason for this ineffability is not that the ideas to be expressed are too high, too spiritual, or too anything-else, but that the forms of feeling and the forms of discursive expression are logically incommensurate, so that any exact concepts of feeling and emotion cannot be projected into the logical form of literal language. Verbal statement [...] is almost useless for conveying knowledge about the precise character of our affective life.”⁸⁷

⁸⁵ Wittgenstein 1953: § 610

⁸⁶ Dewey 1934: 74

⁸⁷ Langer 1957: 91 (her emphasis)

“Certain works of high imagination bewitch the understanding [...], provoking a response which consists in venturing a coherent interpretation of the work which we know in advance to be inadequate, for we also know the value of the interpretation lies in the fact that it escapes us.”⁸⁸

John Dewey seems to infer the existence of aesthetic content from the existence of art itself; Suzanne Langer claims that art expresses feelings and emotions more aptly than language; Arthur Danto seems to argue that aesthetic interpretation is valuable because of the impossibility of providing a coherent interpretation of the artwork. I think it is fair to say that none of these three accounts can be used as a satisfactory account of aesthetic content: Dewey’s inference from the existence of art to the existence of a specific kind of aesthetic content is invalid; Langer’s argument that language cannot express feelings and emotions as precisely as art is unmotivated; and Danto’s attempt to link the value of art (its interpretation, that is) with an alleged impossibility of attaining coherence is at best confusing and at worst confused.

Perhaps it is wrong to expect that there can be an *argument* for the existence of aesthetic content (after all, we hardly ever ask religious people to provide an argument to the effect that they have had an experience with religious content); maybe we ought to be satisfied with the fact that many people *report* that their aesthetic experience involved ineffable content. The aesthetician as I portray him takes experiences to be able to communicate two kinds of content: representational content and aesthetic, non-representational content. The aesthetician holds that, even if the content of a work of art is partly representational (think of a still-life or a photograph), there is a non-representational part as well. It is this non-representational content which accounts for the ineffability felt in aesthetic experiences. The anti-aesthetician denies this.

Note that there are at least two ways for the anti-aesthetician to deny the aesthetician’s view. Either, he could argue that aesthetic content is always representational and therefore, always effable. Or, he could reject the claim that all experiences communicate some kind of content and argue that some are contentless. In this case, the ineffability of some experiences is not due to there being some kind of aesthetic content but due to a complete absence of content, as in the uninteresting cases of the stone or the taste of saffron. The first line of attack is, I think, easily refuted. There are numerous examples of abstract art or pure music which refute the claim that the content of works of art is always representational (Karlheinz Stockhausen’s ‘serial compositions’ come to mind). So if the anti-aesthetician wants to refute the aesthetician’s

⁸⁸ Danto 1973: 45

claim, he has to take the second line. This is what William Kennick does.

5.2.2.1 Aesthetic Experience – Trivially Ineffable?

In ‘Art and the Ineffable’, Kennick argues that we cannot “experience” a proposition, but we can experience works of art. This is because, as he explains, “poems and paintings have meanings and values that ordinary prose descriptions do not have.”⁸⁹ Although his conflated use of the words ‘meaning’ and ‘value’ seems slightly confusing in this context (given that we tend to associate both with propositional content, which is exactly what he opposes them to in this passage), he clarifies what he means in the next sentence:

“Pictures have colour, depth, balance, chiaroscuro, design; propositions do not. Poems have measured rhythm and rhyme, striking images and involved metaphors; ordinary prose statements do not.”⁹⁰

So what Kennick is probably referring to are the specific perceptual qualities of artworks: the colour of a painting stimulates our visual perception, the rhythm of a poem stimulates our auditory perception, etc. He seems to think that propositions don’t have such perceptual qualities, i.e. they don’t stimulate our perceptual apparatus because they are purely cognitive entities. On this picture, the claim that the content of an artwork is ineffable because it cannot be translated into literal language is misguided. He illuminates his position with an analogy:

“Suppose a man complained that there must be something wrong with his protractor because with it he could not draw a square circle. This complaint is unreasonable at best and at worst no complaint at all. The complaints against language that we have been considering are, I think, of this order. They do not strike at remediable difficulties, and hence they do not strike at inadequacies or defects of language at all; they do not point to something that language cannot do or that art can do better than language can. It is one thing to complain that one does not have the words, or the right words, to express or to describe his feelings; it is quite another thing to complain that there are and can be no words to express or describe his feelings.”⁹¹

Kennick thus holds that there is no reason to suppose that an artwork’s “informative content”,⁹² if there is any, cannot be rendered in literal language. His view seems to be that, by calling the specific perceptual qualities of a work of art ineffable, we are making a category mistake. The perceptual qualities of a work of art are just not the kind of thing that could be expressed linguistically (namely, a truth-apt, representational piece of

⁸⁹ Kennick 1961: 316

⁹⁰ Kennick 1961: 316

⁹¹ Kennick 1961: 318

⁹² Kennick 1961: 320

content). The reason why we feel that works of art cannot be “paraphrased” is not because they communicate a specific kind of aesthetic content, but because they afford a specific kind of perceptual experience which lacks content altogether.⁹³

In a related vein, Stephen Davies argues against the idea that the ineffability that attaches to artworks is in any way significant. He writes:

“[...] there is a temptation to develop the view that we have ineffable, perceptual knowledge into the claim that there are ineffable truths that cannot be expressed in language and, further, that such truths are somehow more important or vital than those language can capture. To the extent that they present vital truths through our experience of them, artworks will be thought to have something important to 'say' that cannot be put literally into words. This temptation should be resisted... what is known ineffably is a degree of detail and resolution that in auditory expressions that exceeds the possibility of verbal description or specification, and not some new dimension of the eternal verities.”⁹⁴

Davies thus explains aesthetic ineffability in terms of a lack of descriptive resources to capture a perceptual experience to the full extent. However, since this kind of ineffability is not restricted to aesthetic contexts but can be found in ordinary perception as well, he concludes that we have no reason to consider aesthetic ineffability any more interesting than ordinary perceptual ineffability. The view that the ineffability attributed to artworks is simply a matter of a lack of conceptual nuances is also defended by Diana Raffman, who explores the relationship between human memory capabilities and verbalization within the context of tonal music.⁹⁵ However, I submit that both Raffman and Davies fail to account for the feeling of *significance* or *importance* that is usually reported with reference to ineffable aesthetic experiences, but not with regard to ordinary perceptual experiences. It seems to me that, even if we had the conceptual resources to express all nuances of an artwork (say, the exact colour and position of a painting's brush strokes), we would still not feel that we had managed to express the specifically aesthetic content. Thus, their accounts fail to account for the most interesting aspect of aesthetic ineffability.

⁹³ This view is also held by Malcolm Budd; 1996: 84f.

⁹⁴ Davies 1994: 161

⁹⁵ Cf. Raffman 1993

5.2.2.2 A Twofold Structure of Awareness?

In ‘Aesthetic Ineffability’, Rafael De Clercq applies Michael Polanyi’s⁹⁶ account of the twofold structure of awareness to aesthetic perception. He begins with the following statement:

“Discussing art clearly requires a range of (special) skills. But no matter how perceptive we are, and no matter how well we succeed in communicating our findings to others, there is some point at which even our best efforts run up against the limits of language. That is, at some point we must recognize that much of what we find of significance in art, and in aesthetic objects in general, cannot be rendered in words (without remainder) and so can never become fully our own. This observation, which I take to be in line with common sense, could also be phrased as follows: language, at least in its literal mode, is not able to capture fully the content of an aesthetic experience; aesthetic experience, therefore, may be said to put us in touch with the unsayable or ‘ineffable’.”⁹⁷

De Clercq identifies three questions any account of aesthetic ineffability should be able to answer. The first one is why the content of an aesthetic experience cannot be fully articulated. The second is why we would experience this inarticulable content as ineffable. The third is why we attach importance to what we fail to express in aesthetic experience.⁹⁸ He credits Polanyi’s account of awareness to be capable of answering the first two of these questions. According to Polanyi, perceptual attention (or awareness; he uses them synonymously) has a twofold structure: there is focal awareness and there is subsidiary awareness. A perceptual experience consists of sensory elements, some of which we attend to focally and some of which we attend to subsidiarily. If I understand De Clercq correctly, the way he thinks we attend to an artwork like Beethoven’s *Moonlight Sonata* is by focussing on certain elements (the key, the rhythm, etc.), while leaving other, subsidiary elements in the background. The subsidiary elements of an aesthetic experience are unspecifiable (and thus ineffable) because “as soon as we shift our attention to it, and start to examine it focally, its meaning changes, that meaning being, in fact, the focus upon which it used to bear.”⁹⁹ In other words, as soon as we try to focus on the subsidiary elements, they stop being subsidiary and thus, change their character. I take it that De Clercq would argue that this is what accounts for the often reported feeling of ‘elusiveness’ of aesthetic experiences. Consequently, it is impossible to attend

⁹⁶ Polanyi 1975

⁹⁷ De Clercq 2000: 87

⁹⁸ De Clercq 2000: 87

⁹⁹ De Clercq 2000: 94

to an object both focally and subsidiarily; focal and subsidiary attention are mutually exclusive. Aesthetic experience, on this account, can now be explained as follows:

“[O]ne could say that what is ineffable about a musical work or any other aesthetic object is not some particular part of it to which we are attending with subsidiary awareness (the bass line, the melody, the brush strokes, the rhyme, etc.). It is rather the way these various parts contribute to and are integrated into a whole that can move us deeply.”¹⁰⁰

I think De Clercq’s account is very appealing because it offers a clear explanation of what exactly it is about an artwork that causes the feeling of ineffability. However, despite its merits, the account is inadequate because nothing in De Clercq’s description of aesthetic experiences is unique to *aesthetic* experiences: if Polanyi’s description of a twofold structure of attention is correct, it arguably applies to ‘ordinary’ perceptual experiences as well. Consequently, De Clercq’s account fails to explain why it is that we credit only aesthetic (but not ordinary perceptual) experiences with an extraordinary importance.

5.2.2.3 Non-inferential Content?

In ‘Non-conceptual Content and the Sound of Music’,¹⁰¹ Michael Luntley makes a case for the existence of non-conceptual content by arguing that sometimes experience features content that does not function in the way conceptual content does. His account is designed to evade what he calls McDowell’s ‘kidnapping strategy’ (see 5.2.1 above). Luntley argues that the content of a person’s perceptual experience can only be ‘kidnapped’ by a demonstrative concept if it has the potential to figure in a person’s ‘space of reasons’, i.e. if it has the potential to be employed in rational inference informing a person’s reasons for action. According to Luntley, the distinguishing mark of conceptual content is that it figures in a subject’s reasons which are deployed for belief and action. Nonconceptual content, though representational, cannot perform this function. He writes:

“If subject *S* has a conceptual capacity for discriminating *F*-ness (their experience represents *F*-ness in some way), the representation of *F*-ness must be capable of contributing to the rational organization of their behaviour by figuring in their inferential reasons for belief/action. [...] The capacity to discriminate *F*-ness as contributing to the content of the subject’s experience is nonconceptual if and only if the capacity cannot contribute to the subject’s rational organisation of their behaviour.”¹⁰²

¹⁰⁰ De Clercq 2000: 96

¹⁰¹ Cf. Luntley 2003

¹⁰² Luntley 2003: 406/407

Luntley finds examples of nonconceptual content as he describes it in musical experiences, for example:

“I suggest that the sense of tonality produced by hearing the dominant 7th is, for the novice, a nonconceptual representational content.”¹⁰³

I think that Luntley’s argument for nonconceptual content holds. If aesthetic experience is a sub-class of experience, it makes sense to explain aesthetic experience in comparison to other kinds of experience. Many of our experiences serve to provide information which we will then use in order to rationally organize our thought and behaviour. I will call this kind of experience, which is primarily aimed at the rational organization of behaviour ‘rational experience’. ‘Aesthetic experience’ can then be defined as the kind of experience which is not primarily aimed at the rational organization of behaviour. However, I think Luntley is wrong to insist that it is representational content. Rather, he should have argued that it is a certain type of knowledge. When he claims that the kind of experience he describes is “primarily aesthetic” and that it is “an experience that creates a certain impression, but it is not primarily an experience productive of a rational response”,¹⁰⁴ he describes what is essential about that kind of experience: it is not mirrored in the subject’s actions. For example, experiencing an abstract painting or a piece of pure music as “acquiring a point”¹⁰⁵ involves, according to Luntley, having an expectation about “the balance between elements”¹⁰⁶ of the artwork. And having such an expectation is identical to

“having a representational content without, as it were, knowing what to do with it. The subject of experience with non-conceptual content cannot deploy this content in rational action.”¹⁰⁷

I suggest that the experience Luntley describes is not a non-rational representational experience but the experience of acquiring ineffable knowledge (more on this in 5.3).

To conclude, none of the accounts I examined manages to provide a good explanation of the ineffability of aesthetic experiences. The reason for this common failure is, I maintain, that all these accounts are trying to explain aesthetic ineffability in terms of a specific kind of content. Such attempts are all bound to fail. In the next section, I will argue that what is experienced as ineffable in aesthetic experiences must be explained in terms of ineffable knowledge.

¹⁰³ Luntley 2003: 415

¹⁰⁴ Luntley 2003: 421

¹⁰⁵ Luntley 2003: 422

¹⁰⁶ Luntley 2003: 422

¹⁰⁷ Luntley 2003: 423f

5.3 Ineffable Knowledge?

I have argued in the preceding sections that we cannot explain ineffability in terms of ineffable truths or ineffable content. Both notions either turn out to be incoherent or face decisive objections. If ineffability is thus not a matter of ineffable truths or ineffable contents, what else could it be? In the following section, I will suggest that ineffability must be explained in terms of ineffable knowledge. A preliminary aspect to note is that ineffable knowledge must not be confused with knowledge of the unknowable (which would be a contradiction in terms). The following quote helps to get clear about the distinction:

“The ineffable must not be confused with the unknowable. One may assert [...] that some things [like Kant’s things-in-themselves] are both ineffable and unknowable; nevertheless, the concept of the ineffable is usually so construed that it makes sense to say, ‘I know something somehow, but I cannot put it into words; I cannot say what I know.’”¹⁰⁸

I will consider two principle candidates for ineffable knowledge: standard examples of knowledge-how, and what I call ‘Moorean’ ineffable knowledge.¹⁰⁹ I will critically discuss some recent accounts of knowledge-how and argue that they fail to account for that aspect which we consider ineffable. Following that, I will outline A.W. Moore’s account of ineffable knowledge. I shall argue that his account can be expanded in such a way that all philosophically interesting cases of ineffability can be explained in terms of Moorean ineffable knowledge.

5.3.1 Knowledge-how

The knowledge-how/knowledge-that debate is directly relevant to my discussion of ineffability. The proponents of knowledge-how argue that there is a fundamental difference between knowledge-how and knowledge-that. The alleged difference is that knowledge-that is *propositional* knowledge whereas knowledge-how is *non-propositional*. If it could be shown that some sorts of knowledge are irreducible to propositional knowledge, then arguably some sorts of knowledge are ineffable. Jason Stanley and Timothy Williamson reject the idea that knowledge-how and knowledge-that are different species of knowledge.¹¹⁰ Rather, they argue that knowledge-how is nothing but a particular sub-class of knowledge-that.

¹⁰⁸ Kennick 1967: 181

¹⁰⁹ By which I am referring to A.W. Moore’s account of ineffable knowledge.

¹¹⁰ Stanley/Williamson 2001

The *locus classicus* of the knowledge-how/knowledge-that debate can be found in chapter two of Ryle's *The Concept of Mind*. Ryle argues that, while knowledge-that is clearly propositional knowledge, knowledge-how is non-propositional knowledge which is best described as an ability, a skill, a competence, or a capacity. He writes:

“[Opponents of knowledge-how] are apt to try to reassimilate knowing *how* to knowing *that* by arguing that intelligent performance involves the observance of rules, or the application of criteria. It follows that the operation which is characterized as intelligent must be preceded by an intellectual acknowledgement of these rules or criteria; that is, the agent must first go through the internal process of avowing to himself certain propositions about what is to be done [...].”¹¹¹

Ryle rejects this view because he detects a vicious regress:

“The consideration of propositions is itself an operation the execution of which can be more or less intelligent, less or more stupid. But if, for any operation to be intelligently executed, a prior theoretical operation had first to be performed and performed intelligently, it would be a logical impossibility for anyone ever to break into the circle.”¹¹²

In other words, Ryle holds that it is absurd to assume that engaging in an intelligent action¹¹³ or performance requires the consideration of an infinite number of propositions. His argument is a *reductio ad absurdum*. The assumption for reductio is that knowledge how to ϕ is knowledge that $p(\phi)$. Assuming that considering a proposition is a kind of action, the argument goes something like this:

- a) If a person S engages in an intentional action ϕ , then S knows how to ϕ .
- b) Knowing how to ϕ involves considering the proposition $p(\phi)$.
- c) If S considers the proposition $p(\phi)$, S engages in an intentional action; she knows how to consider $p(\phi)$.
- d) Knowing how to consider $p(\phi)$ involves considering the proposition $q(p(\phi))$.
- e) If S considers the proposition $q(p(\phi))$, S engages in an intentional action; she knows how to consider $q(p(\phi))$. Etc.

Ryle's argument rests on two premises:

- 1) For any intentional action ϕ , by ϕ -ing one employs knowledge *how* to ϕ .
- 2) Employing knowledge *that* p , one considers the proposition that p .

¹¹¹ Ryle 1949: 29

¹¹² Ryle 1949: 30

¹¹³ By an 'intelligent' action, I will understand an 'intentional' action. Unintentional "actions" like developing a rash or breaking one's leg while skiing are arguably not even proper actions, or at any rate, are not intelligent actions. Stanley/Williamson apply the same reading.

Stanley/Williamson argue that Ryle's argument is unsound because there is no uniform reading on which both premises turn out true. Given that we perform many actions without intentionally considering a corresponding proposition, premise (2) can only be true if 'considering a proposition' is interpreted as some kind of unintentional action. However, premise (1) can only be true if it is restricted to intentional actions. Therefore, Stanley/Williamson conclude, Ryle's argument for the fundamental difference between knowledge-*how* and knowledge-*that* cannot get off the ground. After their refutation of Ryle's argument, Stanley/Williamson proceed to an additional linguistic argument (which cannot be discussed in detail here) to the effect that ascriptions of knowledge-*that* and ascriptions of knowledge-*how* neither differ with regard to their syntactic structure nor with regard to their semantics.

They then present their own positive account of knowledge-*how* which can be summarized as follows: for any proposition *r* ascribing knowledge-*how* to a person *S* – let's use the proposition 'Shula knows how to strike the right note' for example – the following holds:

(SW¹¹⁴) *r* is true iff, for some contextually relevant way *w* which is a way for Shula to strike the right note, there is a practical mode of presentation *m*, such that Shula knows under *m* that *w* is a way for her to strike the right note.¹¹⁵

Before entering the discussion of their account, I would like to get the following general worry about formulating knowledge-*how* in terms of knowledge-*that* out of the way.

David Wiggins makes the following remark:

"[A]s Ryle says, 'however many strata of knowing are postulated' and however we imagine the prolonged demonstration itself being bedded down and supplied to the position occupied by '*w*' in '*w* is a way to *V*' (the proposition which is supposedly grasped by one whose competence precisely consist in their 'entertaining' it) 'the same crux always recurs that a fool might have all that knowledge without knowing how to perform, and a sensible or cunning person might know how to perform who had not been introduced to the postulated [proposition that '*w* is a way to *V*'].¹¹⁶

Thus, the worry is that ascribing propositional knowledge (knowledge-*that*) to a person who knows how to ϕ (or *V*, for that matter), seems wrong in cases where the person would herself not be able to formulate that proposition. However, I think this worry is unjustified. A person doesn't need to have been introduced to a proposition in order to know it (at least in cases of propositions ascribing practical knowledge); sometimes one

¹¹⁴ For (Stanley/Williamson)

¹¹⁵ Stanley/Williamson 2001: 430

¹¹⁶ Wiggins 2009: 273

“just knows” it. So there is nothing wrong *per se* with the attempt to express knowledge-how in terms of knowledge-that.

So let’s go back to examine the Stanley/Williamson account of knowledge-how. Clearly, the practical mode of presentation *m* is the most interesting part of this account. It is supposed to capture that which is peculiar about knowledge-how. When we say ‘Shula knows how to strike the right note’, we don’t mean that Shula has watched several choir singers for some time, listened to their explanations about pitch, and now knows certain facts about how to strike the right note. What we mean is that she herself, when singing, is able to strike the right note. This is what Stanley/Williamson intend to capture in the clause ‘there is a practical mode of presentation *m*’ under which Shula knows how to strike the right note. I think this must be understood as follows. In order to grasp *w* as the way to ϕ , a subject *S* must be in a specific receptive state of mind which allows her to identify *w* as the way to ϕ . This receptive state is the practical mode of presentation *m* under which *S* grasps that *w* is a way to ϕ , and under which *S* can demonstrate (*that* is a way to ϕ) and communicate (*is* is a way to ϕ) her knowledge how to ϕ . The role which Stanley/Williamson assign to *m* thus seems to be that of enabling the demonstration and communication of that aspect of knowledge-how which is otherwise not formulatable in language, i.e. the ineffable aspect. Williamson even goes so far as to claim that “[n]o verbal description of [...] *w* is needed; one’s grasp of the propositional content may be distinctively practical (ϕ now!; ϕ like this!)”.¹¹⁷

Unfortunately, however, they don’t provide any further elucidation of *m*. In fact, they even concede that providing a nontrivial characterization of a practical mode of presentation is very difficult.¹¹⁸ They are not concerned by the worry that an analysis of these practical modes of presentation might require appeal to a knowledge-how primitive, since they are “not engaged in the reductive project of reducing talk of knowledge-how to talk that does not involve knowledge-how.” Rather, their only aim is to show “that knowledge-how is a species of knowledge-that.”¹¹⁹ However, they are willing to pay a high price for a relatively modest goal, namely the price of introducing to their account the concept of a practical mode of presentation as a primitive. By doing that, whether or not one concludes that they achieve their goal of translating all knowledge-how into knowledge-that ascriptions, their account fails to explain the most puzzling aspect of knowledge-how. The only indication Stanley/Williamson provide as to

¹¹⁷ Williamson 1997: 44

¹¹⁸ Stanley/Williamson 2001: 429

¹¹⁹ Stanley/Williamson 2001: 433f

what they mean by ‘knowing something under a practical mode of presentation’ is that it “undoubtedly entails the possession of certain complex dispositions”.¹²⁰ However, what else would these dispositions be if not those which ascribe to a person *S* the ability/skill/competence to ϕ ? Yet if it is granted that this is what these dispositions must come down to, then Stanley/Williamson’s account of knowledge-how suddenly becomes very similar again to Ryle’s account of knowledge-how. As Tobias Rosefeldt observes:

“All that Stanley’s and Williamson’s arguments show is that expressions of the form ‘*a* knows how to *F*’ are ambiguous between a reading in which we ascribe some propositional knowledge and one in which we want to say that *a* has a certain ability.”¹²¹

I would go even one step further: Stanley/Williamson’s account does not highlight an *ambiguity* between a propositional reading and an ability reading of knowledge-how ascriptions; rather, it simply highlights that *every* ascription of knowledge-how contains *both* aspects. However, from the fact that knowledge-how ascriptions contain both a propositional aspect and an ability aspect it does not follow that knowledge-how ascriptions can be subsumed under knowledge-that ascriptions. Moreover, given that the irreducible aspect of knowledge-how ascriptions remains an integral part of knowledge-that ascriptions on their account, Stanley/Williamson fail to ‘explain away’ what strikes us as ineffable in knowledge-how ascriptions. The consequence of this is, as Wiggins points out, that knowledge-how reappears as a precondition of identifying the relevant way *w* as a way to ϕ :

“Once we appreciate the difficulty of supposing that knowing the way to *V* (knowing it practically, that is) can be fully accounted for without knowing-how-to-*V* itself reappearing as a precondition for *w*’s being identified as the *w* that verifies the proposition ‘*w* is a/the way by which one *V*’s – once we appreciate everything that would be involved in this – examples [for the reappearance of the ineffable aspect of knowledge-how] will multiply.”¹²²

What Wiggins points out here is that, for the practical mode of presentation to have the desired effect of demonstrating and communicating the relevant piece of knowledge, there must be some sort of incipient knowledge-how in the hearer. For someone to grasp the meaning of the sentence ‘*w* is a way to ϕ ’, it is necessary that that person receive the information under a practical mode of presentation; that, in turn, requires her to be in the relevant receptive state of mind which allows her to receive the information that *w* is a way to ϕ , which then allows her to understand that *w* is a way to ϕ . To put it bluntly: a

¹²⁰ Stanley/Williamson 2001: 429

¹²¹ Rosefeldt 2004: 378

¹²² Wiggins 2009: 271

person won't come to know that w is a way to ϕ unless she already knows that w is a way to ϕ . I won't understand that w (*that* way) is a way to strike the right note unless I already know how to strike the right note. That is just what it means to know how to do something.

To conclude, Stanley/Williamson's account of knowledge-how seems to succeed in reformulating knowledge-how ascriptions like 'knowing how to ride a bicycle' or 'knowing how to play the violin' into knowledge-that ascriptions at the cost of adding a philosophical primitive, the 'practical mode of presentation', to their account. I have argued that the practical mode of presentation is the most interesting aspect of such knowledge-that ascriptions because it is the ineffable aspect; yet Stanley/Williamson fail to shed further light on this aspect. Even if this is taken to be unproblematic in the context of 'ordinary' knowledge-how ascriptions like 'knowing how to play the violin', 'knowing how to ride a bicycle', etc. (and I remain doubtful that it is entirely unproblematic for the reasons given above), there is a kind of knowledge (which I will elucidate in the following section), which seems to be knowledge-how, even though we cannot say what the correct knowledge-how ascription is. Yet if we cannot say what the correct knowledge-how ascription is, we also cannot reformulate it into a knowledge-that ascription. Therefore, despite the objections it raises, the Stanley/Williamson account does not rule out knowledge-how as a candidate for ineffable knowledge.¹²³

5.3.2 Moorean Ineffable Knowledge

In this final section, I will introduce A.W. Moore's account of ineffable knowledge, discuss some of the objections that have been brought forward against it, and argue that ineffability must be explained in terms of Moorean ineffable knowledge.

5.3.2.1 Knowing Through Being Shown: Moore's account

Moore begins his argument for the existence of ineffable knowledge by asking how one of the most fundamental philosophical paradoxes can be resolved. The paradox, which I will henceforth refer to as the 'Paradox of the Idealistic Appeal' (PIA), arises as follows.

¹²³ An alternative to arguing that ineffable knowledge is a kind of knowledge-how would be to argue that it is a specific kind of 'understanding', 'ineffable understanding'. The term 'understanding' is used in slightly different locations as the term 'knowing': even though we can understand *that* something is the case, or *how* something works, we can also simply understand *something*. The verb 'to understand' does not require an object. In fact, one can detect such a shift in Moore's recent writings as well; while he argues for 'ineffable knowledge' in his (1997), he now seems to prefer referring to 'ineffable understanding' rather than 'ineffable knowledge' (cf. Moore 2003b).

Let's say that detached thought about the world means forming absolute representations of the world. Absolute representations are representations which are from no point of view. A point of view can be defined as a "location in the broadest possible sense".¹²⁴ All the features which define the location of the subject count towards the characterization of the point of view:

"Hence points of view include points in space, points in time, frames of reference, historical and cultural contexts, different roles in personal relationships, points of involvement of other kinds, and the sensory apparatuses of different species."¹²⁵

A point of view is thus a vantage point from which to "look" at the world, as it were, which takes into account certain features of reality and leaves out others. Now let's say that one representation weakly entails another if and only if the latter is a consequence of the former. Let's further say that to endorse a representation indirectly is to produce a representation that weakly entails it. Then the following argument shows that detached thought about the world is possible:

"If reality is something substantial that representations answer to, the same reality in every case, then not only must it be possible to provide an account of the kind just described for any possible true representation, but the part of this account that is used for the indirect endorsement of the representation must be combinable with every other such part into a single conception of reality – call it *C*. [...] Now consider any possible true representation *p* from any point of view π . One of the members of *C* must be derived from the account of how *p* is made true by reality. This account, since it serves for pitting *p* against any other possible true representation, including any possible true representation from a point of view incompatible with π , cannot itself be from π . So given that all the members of *C* are from the same points of view, none of them can be from π . But π was chosen arbitrarily. So none of the members of *C* can be from any point of view. Absolute representations are possible."¹²⁶

For many people, providing an argument to that effect will seem rather superfluous, given that science, especially physics, seems to demonstrate quite impressively what the argument intends to establish. However, since not everyone shares these intuitions about science, it is important to show that there is an independent argument for the possibility of detached thought about the world.

So it seems that, on the assumption that the representations we form of the world are representations of a mind-independent reality, we have good reason to believe that thought about the world that is not dependent on a specific perspective or point of view

¹²⁴ Moore 1997: 6

¹²⁵ Moore 1997: 6

¹²⁶ Moore 1997: 72ff

is possible; and indeed, science seems to fit the bill. However, there is also a strong sense in which thought about the world seems “soaked” (Moore) in perspective: doesn’t forming a representation of the world depend on there being a mind in the first place? And if so, isn’t it possible that the world we form representations of depends for its existence on our forming representations of it? This worry is the underlying motive of many formulations of idealism. The ‘Paradox of the Idealistic Appeal’ can thus be stated as follows:

(PIA)

1. Given that detached thought about the world is possible, we have good reasons to believe that the world exists independently of us.¹²⁷
2. The way in which we experience the world gives us reasons to believe that the world exists dependently on us.
3. (1) and (2) are mutually exclusive; they cannot both be true.
4. Thus, we either need to find a way to decide which reasons have more weight, or we need to find a way to reconcile these two claims.

A.W. Moore’s account of ineffable knowledge can be understood as an attempt to solve PIA by reconciling (1) and (2). His argument can be summarized as follows.¹²⁸

Idealism can be generally defined as the view that the form of the world of which we produce representations depends (at least partly) on our representations. Putting it in more flashy terms, idealism is the view that there is a dependence relation between reality and mind. There are at least two kinds of idealism which can be characterized, very roughly, as follows.

According to what could be referred to as ‘empirical idealism’, the dependence relation between mind and reality is *immanent*, i.e. it somehow belongs to the furniture of the world. There are, of course, numerous ways in which it could be spelled out what exactly it means for a dependence relation to be immanent. For example, one way in which the dependence relation could be pinned down within the world is to claim that there are facts which accommodate the perspectivity contained in the dependence relation. And this, in turn, could be done by assuming the existence of something like perspectival facts (see the arguments by Nagel and Jackson discussed in section 5.1.1). As I have argued above, however, this view is untenable.

¹²⁷ ‘Independently of us’ ought to be spelled out in a suitable way (e.g. independent of our minds, of our thought about the world, of us conceiving the world, etc). However, for the purpose of my argument, it will suffice to assume that the idea has enough intuitive force to be intelligible without an elaborate argument.

¹²⁸ Moore expounds his argument most elaborately in his (1997), chapters 7-9.

According to what could be referred to as ‘transcendental idealism’, the dependence relation between mind and reality is *transcendent*, i.e. it is not part of the furniture of the world. If it is not located in the world, it must be located in our thought about the world. The most prominent way of spelling out what exactly transcendental idealism could come down to has, of course, been provided by Kant (he actually invented the term), who argues in the First Critique that the way in which the human cognitive apparatus functions accounts for the experienced perspectivity of the objects of perception. At the beginning of the ‘Transcendental Aesthetic’ in the *Critique of Pure Reason*, Kant states:

“In whatever way and through whatever means a cognition may relate to objects, that through which it relates immediately to them, and at which all thought as a means is directed as an end, is intuition. This, however, takes place only insofar as the object is given to us; but this in turn, is possible only if it affects the mind in a certain way. The capacity (receptivity) to acquire representations through the way in which we are affected by objects is called sensibility. Objects are therefore given to us by means of sensibility, and it alone affords us intuitions; but they are **thought** through the understanding, and from it arise concepts. But all thought, whether straightaway (*directe*) or through a detour (*indirecte*), must ultimately be related to intuitions, thus, in our case, to sensibility, since there is no other way in which objects can be given to us.”¹²⁹

This passage illuminates the view which lies at the heart of Kant’s transcendental idealism: since we can only relate to objects via *intuition*, the objects in the external world (the “thing-in-itself”) can only be objects of our *thought* but not objects of our *knowledge*. The thing in itself can be *thought of* as that which is causing us, i.e. our minds, to have appearances. Yet it cannot be *known* to play that role because we cannot relate to it in an unmediated fashion:

“Space itself, however, together with time, and, with both, all appearances, are **not things**, but rather nothing but representations, and they cannot exist at all outside our mind; and even the inner and sensible intuition of our mind (as an object of consciousness), the determination of which through the succession of different states is represented in time, is not the real self as it exists in itself, or the transcendental subject, but only an appearance of this to us unknown being, which was given to sensibility. [...] Accordingly, the objects of experience are **never** given **in themselves**, but only in experience, and they do not exist at all outside it.”¹³⁰

Hence, we cannot know the mind-independent world “as it is”. Consequently,

¹²⁹ Kant 1998 [1787]: B33

¹³⁰ Kant 1998 [1787]: B520f/A492f (his emphasis)

“[...] what we are talking about is merely an appearance in space and time, neither of which is a determination of things in themselves, but only of our sensibility; hence what is in them (appearances) are not something in itself, but mere representations, which if they are not given to us (in perception) are encountered nowhere at all.”¹³¹

Kant does not mean to say that the thing-in-itself, i.e. the noumenal world, does not exist. Things-in-themselves exist and exercise “a degree of influence on sense”.¹³² What Kant argues for is a combination of ‘empirical realism’, which preserves the reality of ordinary things given to us in experience, with ‘transcendental idealism’, according to which the properties of objects given to us in experience (like their causal powers and their spatial and temporal location) are determined by our minds. They are so determined because two forms of pure intuition, space and time, structure our sensations into the experience of things in space¹³³ and in time¹³⁴. Space and time are not features of the world as it is but they are *a priori* necessary conditions for any human experience whatsoever. Moreover, human understanding imposes ‘categories’ like causality and dependence, unity and plurality, upon the manifold of sensible intuitions (i.e. upon the received sensory content which has been structured by space and time).¹³⁵ Thus, the objects of ordinary experience are *empirically real* but *transcendentally ideal*: they are not to be identified with anything that lies beyond, and thus transcends, the bounds of possible experience.

Transcendental idealism is a way to reconcile the two mutually exclusive claims of PIA by distinguishing between levels: at a transcendent level, the claim that our representations are representations of a mind-independent reality is rejected; at a non-transcendent level it is affirmed. Take the following claim for example:

(SC) ‘The existence of the physical universe depends on the existence of our representations of it.’

The transcendental idealist would determine the truth-value of this sentence as follows:

(TIC)¹³⁶ SC is transcendentally true but non-transcendentally false.

SC is thus true of “our mind-dependent phenomenological world” but false of “the mind-independent physical world”. However, Moore points out a serious problem with

¹³¹ Kant 1998 [1787]: B522/A494

¹³² Kant 1998 [1787]: B208

¹³³ Kant 1998 [1787]: B37ff

¹³⁴ Kant 1998 [1787]: B46ff

¹³⁵ Kant 1998 [1787]: B106

¹³⁶ For ‘Transcendental Idealist Claim’

transcendental idealism: it is self-refuting. TIC constitutes a representation of precisely the kind which transcendental idealism says we cannot produce:

“It is no good treating this claim as transcendentally true but immanently false. The transcendent interpretation does not exist. If it did, it would not be transcendent. If we really cannot produce a representation to the effect that the physical universe depends on the existence of our representations without saying something false, then the physical universe does not depend on the existence of our representations. At a more general level, we cannot represent limits to what we can represent. For if we cannot represent anything beyond those limits, then we cannot represent our not being able to represent anything beyond those limits.”¹³⁷

TIC is not true of the physical world; it is thus a false claim and cannot serve to reconcile the two mutually exclusive claims of PIA. It is also impossible to argue something like ‘transcendental idealism is true of the way in which human beings conceive of the world but it is not true of the world itself’ because that would imply relativism about truth, which I have rejected above. Also, someone might try to argue that whether or not transcendental idealism is true depends on the theory of truth we employ. However, the truth-value of a proposition never depends on the underlying theory of truth. Both on a correspondence and on a coherence view, only one of the following two claims can be true: ‘The objects of the world depend for their existence on my conceiving them’ or ‘It is not the case that the objects of the world depend for their existence on my conceiving them’. Claiming that transcendental idealism is a truth that cannot be stated is also not a way out, for that would be equivalent to claiming that transcendental idealism is an ineffable truth, which is incoherent (cf. section 5.1). We can see from all this that TIC has no explanatory power because, given that it is false, it doesn’t explain anything about the physical world.

However, despite the fact that transcendental idealism is neither plainly true nor ineffably true, the transcendental idealist’s claim is still very appealing. Something just seems right about the notion that what our representations answer to is in some way dependent on those very representations. Moore explains this appeal in terms of the Wittgensteinian concept of ‘being shown’ something. Having ineffable knowledge (or being in a state of ineffable understanding) is thus identical to being shown something. He summarizes his account as follows:

“While we cannot coherently state that transcendental idealism is true, we are *shown* that it is, where ‘*A* is shown that *x*’ is defined as ‘(i) *A* has ineffable knowledge, and (ii) when an attempt is made to put that knowledge into

¹³⁷ Moore 1997: 117

words, the result is: x' . Provided that we can make sense of (i) and (ii), this proposal has the threefold merit of: avoiding self-stultification; being compatible with the incoherence of transcendental idealism; and providing an account of transcendental idealism's appeal."¹³⁸

Moore argues that we can make sense of (i). If all cases of knowledge can be said to be "dispositional states which can serve to explain purposive behaviour", ¹³⁹ then nothing in that description precludes the possibility of there being ineffable knowledge. He also argues that we can make sense of (ii). Even though any ' x ' serving as what could be called a 'pseudo-expression' of an ineffable insight must be considered nonsense, there is nevertheless a sense in which we can expect convergence on the way in which such pseudo-expressions are picked. The reason why we can expect convergence is that human beings share two characteristics: we are all finite beings, and we all aspire to transcend our finitude. And because this is so, we have

"a shared sense of when a piece of nonsense is 'apt' to replace ' x ' in the schema, where 'aptness' is a quasi-aesthetic attribute, such as might occur in poetics."¹⁴⁰

Consequently, some pieces of nonsense will be more apt, and some less, for that attempt. For example: the nonsensical statement 'The existence of the physical universe depends on the existence of our representations of it' is more apt to express my ineffable insight about the way in which I receive the world than the nonsensical statement 'Tautau bärchen Hynksi', or even 'All the world belongs to me'.

Even though the example of transcendental idealism together with PIA, the Paradox of the Idealistic Appeal, is the most vivid way of introducing the notion of ineffable knowledge, Moore's notion does not only apply to transcendental idealism. Rather, it is relevant in all kinds of different contexts. Here are three additional examples for ineffable knowledge:

Having ineffable knowledge means achieving "an inexpressible insight into how one implements, combines, and exploits whatever one knows, an insight into how one works with, and reasons with, whatever one receives. The knowledge one thereby exercises is a kind of understanding. It is knowledge how to process knowledge. It is constituted by various receptive capacities, in particular conceptual capacities."¹⁴¹

And:

¹³⁸ Moore 1997: xii

¹³⁹ Moore 1997: 173; defining knowledge in this way anticipates and avoids the worry that, if knowledge were defined as a *factive* state of mind, one can only know facts, and facts are always potentially expressible.

¹⁴⁰ Moore 1997: xiif

¹⁴¹ Moore 1997: 189 (my emphasis)

“A simple example is my knowing how to exercise the concept of greenness: my knowing what it is for something to be green. This does not consist in my knowing that anything is the case. It is rather a matter of my having the wherewithal to know that various things are the case. For instance, it enables me to know that the leaf I am looking at is green. But it does not itself consist in my knowing that anything is the case because it does not answer to how the world is.”¹⁴²

And:

“Another example of my ineffable knowledge is my knowing how to act out my autobiography: my being able to make narrative sense of my own life. Even if I do this in a way that renders my life desolate or ugly in various ways, I cannot thereby be said to have misrepresented anything. It is much like my knowing how to exercise concepts. Indeed a very important part of it is my knowing how to exercise concepts. Making narrative sense of my life is to a very significant extent a matter of determining what concepts I shall live with: whether or not, for example, to make use of the concept of chivalry, or that of blasphemy, or that of sin.”¹⁴³

Knowing how to exercise concepts, and, more generally, knowing how to use language seems to be a central example for ineffable knowledge. It is also an example which depends less on, as it were, being “trained” in philosophical thought as the example of transcendental idealism requires. Before I will try to demonstrate the fact that Moorean ineffable knowledge can serve to explain all philosophically interesting cases of ineffability, I will address one important objection against this central example for ineffable knowledge, which I call the ‘objection from externalism’.

5.3.2.2 The Objection from Externalism

Moore names as one of the main examples for ineffable knowledge the knowledge enabling one to understand one’s own idiolect:

“My understanding of English is a prime example. I would certainly count that as ineffable, even though it includes large tracts of effable knowledge such as my knowledge that the past tense of a regular English verb is formed by adding ‘ed’ – or, for that matter, my knowledge that the word ‘rabbit’ denotes rabbits, and my knowledge that the word ‘green’ denotes green things. Understanding, of the sort that I have in mind, has nothing to answer to. Of course, I may think that I know what a particular word in English means and be wrong: I may think that the word ‘rabbit’ denotes hares as well as rabbits. If that is the case, then what I understand is strictly speaking not English. But I do still *have* my understanding.”¹⁴⁴

¹⁴² Moore 2003b: 165

¹⁴³ Moore 2003b: 165

¹⁴⁴ Moore 1997: 184

Timothy Williamson raises the following worry about this particular example of ineffable knowledge:

“Moore’s prime example of ineffable knowledge is one’s understanding of one’s own idiolect. He claims that it has nothing to answer to: ‘rabbit’ denotes hares in my idiolect if I think it does. Work by Burge and others undermines this conception of an idiolect. One is fallible even about one’s own idiolect and one’s own concepts.”¹⁴⁵

A work Williamson may be referring to is Burge’s ‘Individualism and the Mental’¹⁴⁶ in which the following scenario is sketched.

(Scenario 1) A person with a pain in his thigh, call him Larry, goes to the doctor and tells him that he has arthritis in his thigh. This is a false belief because, by definition, one cannot have arthritis in one’s thigh but only in one’s joints. We are then asked to imagine the exact same situation in the next possible world, i.e.

“[w]e are to conceive of a situation in which the patient proceeds from birth through the same course of physical events that he actually does, right to and including the time at which he first reports his fear to his doctor. Precisely the same things (nonintentionally described) happen to him.”¹⁴⁷

Larry’s two mental states in the actual and in the counterfactual world differ in content although his physical and nonintentional mental histories remain entirely the same. The only difference between the two worlds is the conventional usage of the term ‘arthritis’:

“However we describe the patient’s attitudes in the counterfactual situation, it will not be with a term or phrase extensionally equivalent with ‘arthritis’. So the patient’s counterfactual-attitudes contents differ from his actual ones.”¹⁴⁸

Burge then draws the following conclusion: “The difference in [Larry’s] mental contents is attributable to differences in his social environment.”¹⁴⁹ He elaborates his point by mentioning several variations of the thought experiment, each relying on different variants of error like misunderstanding or misconception, partial understanding, etc., and argues that his argument applies to all cases in which it is possible to attribute to someone a mental state or event whose content involves a notion that the subject falsely or incompletely understands. The thought experiment can also be run reversely.

(Scenario 2) We are asked to imagine a person who has full and correct understanding of the term ‘arthritis’ in the actual world and tells the doctor (correctly)

¹⁴⁵ Williamson 1999: 44

¹⁴⁶ Burge 1979. Burge’s work is inspired by Hilary Putnam’s ‘Twin Earth’ argument (Putnam 1975) for semantic externalism.

¹⁴⁷ Burge 1979: 600

¹⁴⁸ Burge 1979: 601

¹⁴⁹ Burge 1979: 601

that he has arthritis. In the counterfactual world, however, ‘arthritis’ applies to rheumatic ailments of all sorts. So the patient’s beliefs in the two worlds (and hence his mental states) differ because of the different usage of the term ‘arthritis’. Burge states:

“The reversal of the thought experiment brings home the important point that *even those propositional attitudes not infected by incomplete understanding* depend for their content on social factors that are independent of the individual, asocially and non-intentionally described. For if the social environment had been appropriately different, the contents of those attitudes would have been different.”¹⁵⁰

Burge’s argument suggests that, in addition to being determined by understanding, inference patterns, and other features, the contents of a person’s attitudes depend on communal practice. As the reverse argument intends to establish, this supposedly holds even when a person fully understands the content of her mental attitude. Burge’s argument thus suggests a form of externalism about mental contents. According to this view, whatever a person thinks or says is *determined* by aspects of the world external to the subject’s mind. This is a stronger claim than merely stating that mental states are *caused* by external factors. It is the claim that a person *S*’s given mental state *ms* could not have existed as it does if *S* had not been embedded in the very external environment which determines the content of *ms*. The relations between *S* and its external environment thus make up the identity of *ms*.

However, I think Moore’s account can be defended against Williamson’s worry that the prime examples of Moore’s account of ineffable knowledge aren’t valid examples at all. For one thing, Moore clearly states that

“[the conceptual capacities which constitute ineffable knowledge] have nothing to answer to. What matters is what can be done with them. One’s insight does not therefore consists of receiving anything. It derives from a certain *mode* of reception, a mode of making sense.”¹⁵¹

This passage emphasizes that Moore’s claim that the understanding of one’s own idiolect constitutes an example of ineffable knowledge does not depend on our language use being infallible. It seems to me that Moore could happily admit that we are sometimes mistaken about the contents of our mental states without having to give up his claim that the understanding of one’s own idiolect constitutes an example of ineffable knowledge. However, since someone might still worry that Williamson’s objection constitutes a problem for Moore, I will briefly sketch two possible defences.

¹⁵⁰ Burge 1979: 605 (his emphasis)

¹⁵¹ Moore 1997: 189

Defence (1): Burge's account is only meant to apply to 'oblique occurrences' of expressions. If an expression occurs obliquely, it functions in a content-clause (i.e. in a 'that'-clause referring to mental states or events which provide the content of someone's intentional state, e.g. belief) in such a way that it cannot be substituted with an extensionally equivalent expression without that substitution changing the truth-value of the sentence. A classic illustration for this is the *Water* – *H₂O*-example: if Bertrand thinks that water is not fit to drink, it does not follow that he thinks that H₂O is not fit to drink, even though water and H₂O can be considered extensionally equivalent. Burge claims that in these two cases, the mental states (and thus the thoughts) are distinct even though the expressions have the same extensions. The opposite of an oblique occurrence of a term is a non-oblique occurrence of a term: Burge claims that there would be no difference in the content of Bertrand's thought 'the water in the glass over there is impure' if water was replaced with H₂O. The reason is that the expression 'water', in this context, is not a determining component of Bertrand's thought content. Burge summarizes:

"The crucial point [...] is the assumption that obliquely occurring expressions in content-clauses are a primary means of identifying a person's intentional mental states or events. [...] For any distinct contents, there will be imaginable contexts of attribution in which, even in the loosest, most informal ways of speaking, those contents would be said to describe different mental states or events. [...] Most of the cases we discuss will involve extensional differences between obliquely occurring counterpart expressions in that-clauses. In such cases, it is particularly natural and normal to take different contents as identifying different mental states or events."¹⁵²

Assuming that Burge's externalist account is correct, the first defence is that Moore could simply add a qualification to his examples of ineffable knowledge, thus restricting it to non-oblique occurrences of expressions. As Burge himself claims, it is impossible that a person be mistaken about her non-obliquely occurring expressions. Williamson's worry that the prime examples of Moore's account of ineffable knowledge aren't examples at all would not apply in this case.

Defence (2): Another way of defending Moore's account is to challenge Burge's argument. This could be done as follows.¹⁵³ Let's look again at Burge's first thought-experiment, scenario (1). Larry thinks he has arthritis in his thigh, not knowing that 'arthritis' denotes a disease which only occurs in people's joints. Larry, who expresses his

¹⁵² Burge 1979: 599

¹⁵³ This argument is due to Crane (1991). I am in no way suggesting that this counter-argument is a decisive objection against all kinds of externalism. I merely want to outline one way in which externalism can be challenged.

belief with the sentence ‘I have arthritis in my thigh’, thus arguably believes two things: the first-order belief that (A) he has arthritis in her thigh, and the second-order belief that (B) the sentence ‘I have arthritis in my thigh’ is the right way to express his belief. In order to account for the mistake attributed to the patient in Burge’s example, one of the following three options must hold: either, Larry’s belief (A) is false, or Larry’s belief (B) is false, or both (A) and (B) are false. The fact that (B) is false is clear from the setup of Burge’s thought experiment – Larry has, by stipulation, a false belief about which sentence expresses his belief correctly. What about belief (A)? How do we correctly describe Larry’s belief (A)? There are two options. Either, we attribute to him the concept ‘tharthritis’ which applies both to arthritis and to the disease he has in his thigh. In that case, he would (truly) believe that he has tharthritis in his thigh. Or, we attribute to Larry the public concept ‘arthritis’. In that case, he would (falsely) believe that he has arthritis in his thigh. As a way to decide which belief-attribution is correct, Crane suggests that we look at the patient’s dispositions:

“[...] a problem with Burge's thesis is presented by the fact that a psychologist would not be able to distinguish between the non-verbal behaviour [...] of two Burgean twins in certain experimental situations [...]. This would seem to be so in our case too: up to the time described by the thought experiment, [the patient] has (*ex hypothesi*) all the same dispositions to (non-verbal) behaviour in the actual and the counterfactual situations. This suggests that we should attribute the concept *tharthritis* to him.”¹⁵⁴

In other words, Crane holds that we have no reason to assume that Larry employs different concepts in the actual and in the counterfactual situations because his dispositions to act in certain ways remain exactly the same. Arguably, the patient has the same concept in both scenarios, namely ‘tharthritis’, whereas he lacks the concept ‘arthritis’ in both scenarios. Thus, in both scenarios, the Larry’s belief is that he has tharthritis in his thigh. The crucial difference between the scenarios is not a difference in the content of the patient’s beliefs, but a difference in the correctness of her utterance: ‘I have arthritis in my thigh’ expresses his belief correctly only in one of the two scenarios, namely in the scenario where ‘arthritis’ denotes a disease in joints and thighs. It seems to me that Crane has thus provided a powerful argument against the intuition that different belief ascriptions can be true of physically identical subjects. I conclude that, even though objections from externalism need to be taken seriously by anyone who wants to defend

¹⁵⁴ Crane 1991: 19

Moorean ineffable knowledge by means of the example of our knowledge how to use language, they don't constitute a decisive objection against its existence.

5.3.2.3 Applications of Moorean Ineffable Knowledge

I will now explain why I think that Moorean ineffable knowledge can explain all philosophically interesting cases ineffability. Of the eight exemplary instances of ineffability I gave at the very beginning, I ruled out the first four as philosophically uninteresting because they did not constitute 'genuine' cases of ineffability according my definition ('nescient' and 'impediment' ineffability), or because they fulfilled the conditions of my definition in a trivial way ('trivial' and 'sensational' ineffability). The eighth instance of ineffability I gave was the case of 'Moorean' ineffability of transcendental insights, which has been explained in this chapter. The remaining three cases of ineffability were 'aesthetic', 'mystical', and 'contradiction' ineffability.

I will begin with 'aesthetic' ineffability. Why is it impossible to paraphrase what is conveyed through a painting, a melody, a poem into literal language? Certainly not because aesthetic experiences reveal some kind of ineffable truth – as argued above, the concept of an ineffable truth is incoherent. In section 5.2, I have argued that the reason can also not be that there is a special kind of *content*, 'aesthetic content', which can only be expressed through artworks. Trying to argue that aesthetic ineffability is a case of 'sensational' (and thus, philosophically uninteresting) ineffability fails to account for the feeling of significance or importance we associate with aesthetic ineffability. At the end of section 5.2, I suggested that the ineffability of aesthetic experience is due to the fact that we have an ineffable insight (in fact, I would even add that works of art seem to *further* ineffable insights¹⁵⁵). An insight into what? An insight into how the world we receive makes sense to us, how we as finite beings with a spatiotemporal location relate to what we receive through our senses and process through our cognition. It is an insight into one's situatedness in the world.

The same can be said for instances of 'mystical' ineffability. I cannot express my knowledge of G-d. I am shown that He exists, but I cannot express that. Again, the felt ineffability is not due to some ineffable truth or ineffable content which resists expression. There is no truth or content to express. In fact, the sentence 'G-d exists' does not constitute a truth but a piece of nonsense which I consider an apt attempt to express

¹⁵⁵ N.B.: Both Arthur Schopenhauer (Part III of *The World as Will and Idea*), and, influenced by him, Theodor W. Adorno (*Aesthetic Theory*) argue at length that the experience of art fosters insights into the human condition and situatedness in the world.

my ineffable insight into my situatedness in the world. Saying ‘G-d exists’ is just the best we can do in order to hint at the ineffable knowledge we receive.

What about ‘contradiction’ ineffability, which consists in the impossibility to express certain states of affairs without producing a contradiction? Does it make sense to speak of an ineffable insight into our situatedness in the world in such cases? I believe it does. In fact, I believe that cases of ‘contradiction’ ineffability are a sub-class of Moorean ineffability.¹⁵⁶ The ineffable insight we gain is an insight into the expressive limits of our language. Our language, just as our thought, has boundaries which cannot be crossed, but which we attempt to cross constantly; we constantly try to put ineffable knowledge into words, be it in the context of religion, or aesthetic experience, but also in mathematics (think about the infinite, or limits of iteration) and most systematically in philosophy. That this is so provides an insight into the human condition, an insight into how (and how intently) we try to make sense of the world we receive, and of our place within it.

6 Conclusion: Ineffability is Making Sense

I began my argument by raising a question: why do we attach importance to what we are unable to put into words? I adumbrated that this was because we feel that there is *something* of importance lying behind the linguistic barrier. The goal of the thesis was to examine to what extent this intuition is justified, particularly whether we have reason to think that there are metaphysically substantial entities of some sort which resist expression, or whether one should doubt that there actually *is* anything in the world one intends to express but fails to. I provided a number of examples in which the term ‘ineffability’ is encountered and offered an initial definition of ‘ineffability’, provided with regard to a standard of expressibility and in contrast to ‘indescribability’. The definition I provided was then used to identify different *types* of ineffability, and to separate philosophically interesting from philosophically uninteresting cases of ineffability. I then proceeded to an analysis of the philosophically interesting cases. The central question was what exactly it is that we want to call ineffable in those cases. I considered three candidates: ineffable truths, ineffable content, and ineffable knowledge. I argued that ineffability cannot be explained in terms of ineffable truths because the concept ‘ineffable truths’ is linguistically incoherent and leads to an incoherent picture of reality. Another argument for ineffable truths which invokes excess proposition can be avoided

¹⁵⁶ For a different view about contradictions and their relation to the human situatedness in the world see Priest 2002.

by adjusting one's definition of expression. I then argued that ineffability can also not be explained in terms of ineffable content. I considered several accounts of nonconceptual, ineffable content in perception and argued that none of them succeeds in establishing the existence of nonconceptual content. I also considered arguments for the existence of ineffable aesthetic content and rejected them on the grounds that they failed to make a convincing case for the existence of aesthetic content as a kind of content different from 'ordinary' content. Moreover, they all failed to account for the feeling of importance we attach to what we cannot express. Finally, I argued that ineffability must be explained in terms of Moorean ineffable knowledge. I discussed to what extent arguments for the distinctness of knowledge-how and knowledge-that are relevant to the examination of ineffability and then proceeded to an exposition and defence of Moorean ineffable knowledge. I concluded that the philosophically interesting cases of ineffability must be explained in terms of Moorean ineffable knowledge.

I take myself to have provided a positive answer to the central question of this enquiry: whether our intuition that something of importance is revealed to us in moments of ineffability is justified. Moments of ineffability are moments in which we acquire ineffable knowledge. Ineffable knowledge provides insights into the limits of our cognition. These limits are determined by our situatedness in the world and determine our ways of making sense of it. Ineffable insights are insights into our place in the world.

Where does my enquiry lead from here? I think there are plenty of interesting issues which could be explored on the basis of my results, but one of them is, I think, particularly intriguing. It is an issue that relates to the ways in which we attempt (and fail) to express the ineffable. Whilst it is self-defeating to attempt to express the ineffable, it is perfectly respectable to try to put into words the experience of being shown something which emerges from the failed attempt to express the inexpressible. The nonsense we produce when we try to express our ineffable knowledge can "feel" more or less "apt" as an attempt to express the ineffable. Since nonsense cannot be evaluated according to standards of truth and falsity, it makes sense to assume that the way in which we choose one piece of nonsense over another for attempting to express an ineffable insight is determined by some kind of aesthetic judgement. What exactly this quasi-aesthetic judgement comes down to is a question which is worth being explored in depth on the basis of this enquiry.

7 Bibliography

- Armstrong, David Malet. 1968. *A Materialist Theory of Mind*. London: Routledge and Kegan Paul.
- Armstrong, David Malet. 1983. 'Indeterminism, proximal stimuli, and perception'. *The Behavioural and Brain Sciences* Vol. 6: pp. 64-65.
- Budd, Malcolm. 1996. *Values of Art: Painting, Poetry, and Music*. London: Penguin Books.
- Burge, Tyler. 1979. 'Individualism and the Mental'. *Midwest Studies in Philosophy* Vol. 4, *Studies in Metaphysics* (edited by Peter French, T. Uehling Jr., and H. Wettsten). Minneapolis: University of Minnesota Press.
- Bermúdez, José Luis. 1998. *The Paradox of Self-Consciousness*. Cambridge, MA: The MIT Press.
- Bermúdez, José Luis. 2007. 'What is at stake in the debate about nonconceptual content?' *Philosophical Perspectives* Vol. 21 No. 1: pp. 55-72.
- Carroll, Lewis. 1994 [1872]. *Through the Looking Glass*. London: Penguin Books.
- Carroll, Noel. 2006. 'Aesthetic Experience: A Question of Content'. In: Matthew Kieran (ed.). *Contemporary Debates in Aesthetics and the Philosophy of Art*. Oxford: Blackwell Publishing: pp. 69-97.
- Chisholm, Roderick. 1957. *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press.
- Crane, Tim. 1988. 'The waterfall illusion'. *Analysis* Vol. 48: pp. 142-147.
- Crane, Tim. 1991. 'All the difference in the world'. *The Philosophical Quarterly* Vol. 41: pp. 1-25.
- Danto, Arthur. 1973. 'Language and the Tao: Some Reflections on Ineffability'. *Journal of Chinese Philosophy* Vol. 1: pp. 45-55.
- Davidson, Donald. 1974. 'On the Very Idea of a Conceptual Scheme'. *Proceedings and Addresses of the American Philosophical Association* Vol. 47: pp. 5-20.
- Davies, Stephen. 1994. *Musical Meaning and Expression*. London: Cornell University Press.
- DeClercq, Rafael. 2000. 'Aesthetic Ineffability'. *Journal of Consciousness Studies* Vol. 7 No. 8-9: pp. 87-97.
- Dennett, Daniel. 1988. 'Quining Qualia'. In: A. Marcel and E. Bisiach (eds). *Consciousness in Modern Science*. Oxford: Oxford University Press.
- Dennett, Daniel. 1991. *Consciousness Explained*. Boston: Little, Brown and Company.
- Dewey, John. 1934. *Art as Experience*. London : G. Allen & Unwin.
- Dretske, Fred. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Dretske, Fred. 1983. 'Précis of "Knowledge and the Flow of Information"'. *The Behavioural and Brain Sciences* Vol. 6: pp. 55-63.
- Evans, Gareth. *The Varieties of Reference*. Oxford: Oxford University Press.

- Gendler, Tamar Szabó, and Hawthorne, John (eds.). 2002. *Conceivability and Possibility*. Oxford: Oxford University Press.
- Hawthorne, John, and Oppy, Graham. 1997. 'Minimalism and Truth'. *Noûs* Vol. 31 No. 2: pp. 170-196.
- Heck, Richard. 2000. 'Nonconceptual Content and the Space of Reasons'. *The Philosophical Review* Vol. 109: pp. 483-523.
- Jackson, Frank. 1977. *Perception: A Representative Theory*. Cambridge: CUP.
- Jackson, Frank. 1982. 'Epiphenomenal Qualia'. *The Philosophical Quarterly* Vol. 32 No. 127: pp. 127-136.
- Jackson, Frank. 1986. 'What Mary didn't know'. *The Journal of Philosophy* Vol. 83 No. 5: pp. 291-295.
- Kant, Immanuel. 1998 [1787]. *Critique of Pure Reason*. (translated by Paul Guyer and Allen Wood). Cambridge: Cambridge University Press.
- Kennick, William. 1961. 'Art and the Ineffable'. *The Journal of Philosophy* Vol. 58 No. 12: pp. 309-320.
- Kennick, William. 1967. 'The Ineffable'. In: Paul Edwards: *The Encyclopaedia of Philosophy*. New York: Macmillan: pp. 181-183.
- Langer, Suzanne. 1957. *Problems of Art*. New York: Holiday House.
- Lewis, David. 1986. *On the Plurality of Worlds*. Oxford: Blackwell Publishing.
- Luntley, Michael. 2003. 'Non-conceptual Content and the Sound of Music'. In: *Mind & Language* Vol. 18 No. 4: pp. 402-426.
- McDowell, John. 1994. *Mind and World*. Cambridge, MA: Harvard University Press.
- Mellor, Hugh. 'Crane's Waterfall Illusion'. *Analysis* Vol. 48: pp. 147-150.
- Moore, A.W. 1989. 'A Problem for Intuitionism: The Apparent Possibility of Performing Infinitely Many Tasks in a Finite Time'. *Proceedings of the Aristotelian Society*, New Series, Vol. 90: pp. 17-34.
- Moore, A.W. 1997. *Points of View*. Oxford: Oxford University Press.
- Moore, A.W. 2003a. 'Ineffability and Nonsense'. *Proceedings of the Aristotelian Society*, Supplementary Volumes, Vol. 77: pp. 169-193.
- Moore, A.W. 2003b. 'Ineffability and Religion'. *European Journal of Philosophy* Vol. 11 No. 2: pp. 161-176.
- Nagel, Thomas. 1974. 'What Is It Like to Be a Bat?' *The Philosophical Review* Vol. 83 No. 4: pp. 435-450.
- Nagel, Thomas. 1986. *The View from Nowhere*. Oxford: Oxford University Press.
- Peacocke, Christopher. 1992. *A Study of Concepts*. Cambridge, MA: MIT Press.

- Peacocke, Christopher. 2001a. 'Does Perception have a Nonconceptual Content?'. *The Journal of Philosophy* Vol. 98 No. 5: pp. 239-264.
- Peacocke, Christopher. 2001b. 'Phenomenology and nonconceptual content'. *Philosophy and Phenomenological Research* Vol. 58: pp. 381-388.
- Plato. 1961 [428-348 BCE]. *The collected Dialogues of Plato, including the Letters*, edited by Edith Hamilton and Huntington Cairns. New York: Random House.
- Polanyi, Michael, and Prosch, Harry. 1975. *Meaning*. Chicago: University of Chicago Press.
- Priest, Graham. 2002. *Beyond the Limits of Thought*. Oxford: Oxford University Press.
- Putnam, Hilary. 1975. 'The Meaning of "Meaning"'. *Philosophical Papers Vol. II : Mind, Language, and Reality*. Cambridge: Cambridge University Press.
- Raffman, Diana. 1993. *Language, Music, and Mind*. Cambridge, MA: The MIT Press.
- Reid, Thomas. 2003 [1764]. *An Inquiry into the Human Mind, on the Principles of Common Sense*. Electronic reproduction. Farmington Hills, Mich. : Thomson Gale. Available at: <http://galenet.galegroup.com/servlet/ECCO?c=1&stp=Author&ste=11&af=BN&ae=T110558&tiPG=1&dd=0&dc=fle&docNum=CW117422497&vrsn=1.0&srchtp=a&d4=0.33&n=10&SU=0LRF&locID=oxford>
- Rosefeldt, Tobias. 2004. 'Is Knowing-how Simply a Case of Knowing-that?'. *Philosophical Investigations* Vol. 27 no. 4: pp. 370-379.
- Russell, Bertrand. 1935. 'The Limits of Empiricism'. *Proceedings of the Aristotelian Society* Vol. 36: pp. 131-150.
- Ryle, Gilbert. 1949. *The Concept of Mind*. Chicago: Chicago University Press.
- Savage, Richard. 1972. *Foundations of Statistics*. New York: Dover Publications.
- Siegel, Susanna. 2005. 'The Contents of Perception'. *Stanford Encyclopaedia of Philosophy*. URL: <http://plato.stanford.edu/entries/perception-contents/>
- Stanley, Jason and Williamson, Timothy. 2001. 'Knowing How'. *The Journal of Philosophy* Vol. 98 No. 8: pp. 411-444.
- Thomson, James F. 1969. 'Truth-Bearers and the Trouble about Propositions'. *The Journal of Philosophy* Vol. 66 No. 21: pp. 737-747.
- Travis, Charles. 2004. 'The Silence of the Senses'. *Mind* Vol. 113 No. 449: 57-94.
- Waismann, Friedrich. 1951. *Introduction to Mathematical Thinking. The Formulation of Concepts in Modern Mathematics*. New York: Harper & Row.
- Wiggins, David. 2009. 'Knowing How To and Knowing That'. In: Hans-Johann Glock and John Hyman (eds.). *Wittgenstein and Analytic Philosophy. Essays for P.M.S. Hacker*. Oxford: Oxford University Press: pp. 263-277.
- Williamson, Timothy. 1999. 'Review of A.W. Moore, *Points of View*'. *Philosophical Books* Vol. 40 No. 1: pp. 43-45.