

The Context-Sensitivity of Thought

PhD in Philosophy

Philosophy Department

Neil Hamilton Fairley

September 2019

Declaration: I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Acknowledgements

The production of this thesis would not be possible without the help of support, both academic and personal, for which I am very grateful. On the academic side of things, the patience, prompt advice and feedback from my supervisors Professor Emma Borg and Dr. Nat Hansen was invaluable. Also, of immense help were the PhD students at Reading whose helpful comments on draft chapters greatly improved them. In no particular order, I would like to thank Sami Rissanen, Sarah Fisher, Antonio Scarafone, Marat Жардимгалиев, Aart Van Gils, Lanell Mason, Dr. Jumbly Grindrod, Dr. Michela Bariselli and Dr. Euan Metz. The Philosophy Department at Reading generally provided a great deal of support and opportunities including the regular provision of visiting speakers, fascinating reading groups, teaching positions and conference organisation.

I would also like to thank those who gave me feedback outside of the University of Reading at conferences, workshops and summer schools. In particular those who attended The Norwegian Summer Institute of Language and Mind, Philosophy of Language and Mind Conference 4 and the Second Context Communication Cognition Conference.

I would also like to thank the University of Reading for providing me with a studentship to pursue this research, as well as a position of Hall Mentor that allowed me to continue living in Reading.

Turning to personal support, I would like to thank the PhD students past and present in the Philosophy Department. Also, my former flatmates John, Djamila, Yiming, Ijioma, Sam and Shannon. I would also like to thank Thais, Milena, Katja, Carrie-Anne, Feroz, Louis and Metin. Last but not least I would like to thank my family.

The Context-Sensitivity of Thought

Contents

The Context-Sensitivity of Thought.....	3
Abstract.....	6
1. Introduction	7
1.1 What is Context-Sensitivity?	11
1.1.1 Examples of Context-Sensitivity.....	11
1.1.2 Defining Terms	15
1.1.3 Underdeterminacy	20
1.2 Views of Context-Sensitivity in Natural Language	24
1.2.1 Really Radical Contextualism	25
1.2.2 Radical Contextualism	27
1.2.3 Moderate Contextualism	28
1.2.4 Semantic Minimalism.....	29
1.2.5 No Context-Sensitivity.....	32
1.3 Near-Side vs. Far-Side Pragmatics.....	33
1.3.1 More Examples of Context-Sensitivity	34
1.4 Scope of the Thesis	40
1.4.1 Antithesis.....	40
1.4.2 Thesis.....	42
1.5 Summary	45
1.5.1 Chapter Outlines	45
2. Thoughts.....	50
2.1 The Language of Thought Hypothesis.....	51
2.1.1 The Representational Theory of Mind and The Computational Theory of Mind	54
2.1.1.1 Representational Theory of Mind (RTM)	54
2.1.1.2 The Computational Theory of Mind (CTM)	57
2.2 Modules	78
2.2.1 Massive Modularity.....	82
2.2.2 Heuristics.....	84
2.3 Defence of a Modular Account	86
2.3.1 Cognitive Penetration of Modules	86
2.3.2 Domain Specificity.....	92

2.3.3 No Explanatory Work	94
2.4 Concepts.....	97
2.5 A Sentence in Mentalese	101
2.6 The Mixed View.....	103
2.6.1 Comparisons with the Mixed View	104
2.6.2 Mixed View.....	106
2.6.3 Summary	115
2.7 The Thesis Revisited	116
2.8 Summary	122
3. De Se Attitudes and Computation	124
3.1 Inconsistent Triad.....	126
3.2 Claim (i)	126
3.2.1 Expanded Doctrine of Propositions	132
3.2.2 The Challenge to <i>De Se Essentialism</i>	133
3.2.3 The Problem Posed by <i>De Se Attitudes</i>	135
3.2.4 Defending the Problem of the <i>De Se Attitudes</i>	137
3.3 Claim (ii)	141
3.3.1 Defence of claim (ii)	144
3.3.2 Millikan’s View, the Case Against (ii)	145
3.3.3 Responses to Millikan	146
3.3.4 Indexing <i>De Se Attitudes</i>	149
3.3.5 Over Intellectualising?.....	150
3.4 Claim (iii).....	151
3.4.1 Defence of claim (iii)	154
3.4.2 Responses to Argument 1	159
3.5 Rejecting Premise (1)	163
3.5.1 Restricting the Scope of the CTM.....	164
3.5.2 Summary	169
3.6 Conclusion	170
4. Searles Background Theory of Intentionality: Cause for Concern	172
4.1 Introduction	172
4.2 Searle’s Background Theory.....	173
4.3 Perry’s Concerning Relationship	185
4.4 Objections	190

4.4.1 A Slippery Slope?	190
4.4.2 A Flawed Analogy?	195
4.4.3 Holism	197
4.5 Conclusion	207
5. Travis and Judgement	210
5.1 Fodor	211
5.1.1 Recap	211
5.1.2 Additional Detail	216
5.2 Travis's Criticisms	217
5.2.1 Rule-Following Arguments	222
5.2.2 Token-Underdeterminacy	224
5.3 Travis's view	235
5.3.1 Judgement	238
5.3.2 Precondition 1 – Purpose	250
5.3.3 Precondition 2 – The Parochial sensibility.	255
5.4 The Regress Argument	259
5.4.1 Do Concepts Matter?	261
5.5 Conclusion	267
6. Conclusion	270
Bibliography	278

Abstract

I defend the claim that it is possible for thoughts to be context-sensitive. Assuming that a thought is a sentence of Mentalese (i.e. a well-formed string of mental representations) and content is a function from indices to truth-values, then a thought, T, is context-sensitive IFF at least one of the following three conditions are met:

- (i) T exhibits character-underdeterminacy, where T is character underdetermined iff a component of T makes an explicit reference to the context to establish content.
- (ii) T exhibits type-underdeterminacy, where T is type underdetermined iff there are tokens of T that have distinct truth-values.
- (iii) T is token-underdetermined, where T is token underdetermined iff for some possible states of affairs its truth-value is indeterminate.

Additionally, there must be a mechanism by which the content of T can be determined by appeal to the context. Otherwise T is not sensitive to the context.

The most significant objection to the possibility of context-sensitive thoughts comes from a regress argument. The concern is that any process of removing context-sensitivity from a thought will token an additional thought. If that further thought is context-sensitive, then another thought will have to be tokened. This process will only terminate when a non-context-sensitive thought is arrived at. Yet this seems to show that context-sensitivity in thought, if possible, can only be a contingent, detachable feature of thoughts.

To counter this argument, I will present three versions of context-sensitive thought, mirroring each of (i)-(iii). These views make use of a range of relationships between thoughts and their contents or extensions that are affected by the context, but I will argue they do not require the generation of an additional thought. This provides a range of views on which thoughts can be treated as essentially context-sensitive.

1. Introduction

The thesis that I will be defending is that it is possible to understand sentences that we have in thought as being context-sensitive in the same way that we can understand natural language sentences as being context-sensitive. This is the claim that the proposition associated with a sentence (in thought) can vary with the context in which it is tokened. These sentences will be understood as Mentalese sentences, a notion I explain more fully in chapter 2. For the moment, Mentalese can be understood as having a vocabulary of concepts that can be composed in a manner comparable to natural language words. This approach has the advantage of making thoughts comparable to natural language sentences and makes it possible to consider whether they share other properties. I will claim that foundational sentences in thought could be context-sensitive and will elaborate on what this means below. For a thought to be foundational is for it to be true in virtue of no other thought. This means that if any thoughts are foundational and context-sensitive, that context-sensitivity cannot be removed by making use of a context insensitive thought.

To investigate this claim, I will look at three kinds of context-sensitivity that are thought to occur in natural language sentences, when they are uttered. These models are indexicals (such as “I”, “you”, “here”, “now”), the Background theory of content (according to which representations only express a proposition by making use of a set of “background” assumptions) and a judgement dependent view (according to which objects only fall under a given concept when they are correctly judged to do so). If any of these models of context-sensitivity can exist in thought without contradiction (or a regress or some other a priori worry) that will be a point in favour of the thesis. If these three views are plausible, that will provide even stronger evidence in favour of this thesis.

In opposition to these views is the thesis that there could not be context-sensitivity at the level of thought. This is often referred to as the Mixed View (MV), which maintains (roughly) that sentences

in natural language can be context-sensitive, but that thoughts cannot be.¹ Proponents of this view include Jerry Fodor and possibly Robyn Carston.² This would mean that for any given thought, it will not admit of context-sensitivity. There is also a weaker version of this view, which is that at the foundational level of thought, there cannot be any thoughts that are context-sensitive. A thought is foundational if its truth does not depend on the truth of any other thought.³ On this weaker version, there can be context-sensitive thoughts, but there is no context-sensitivity at the foundational level of thought.

The Mixed View (MV) gets some of its plausibility from apparent difficulties in allowing for the context-sensitivity of thought. To give some of the motivation for this view, one might think that thoughts are composed of concepts and that a given concept type should always have the same content across the contexts in which it is tokened. For some, that is just what it is for a concept to be of the given type.⁴ If that is so, then there will not be any variation in a concept type's content across contexts of use. If there were variation in the concepts' content across contexts of use, then we might think that there are really just different concept types being described under the same name. So, if we have the same concepts in the same order, they'll express the same content on this view.

Another powerful motivator of this view is the claim that once we reach (foundational) thoughts, there is nowhere else to go that can take account of the role the context plays in natural language. Thoughts are the bedrock of representation.⁵ This assumption seems to underlie the idea that there might be a translational semantics for natural languages.⁶ On this approach a natural language sentence has the semantics that it does because it is translated into a thought. This sort of move is taken to account for the ways in which natural languages are sometimes taken to be context-

¹ Jaque 2017: 4

² See Fodor (2001) or Carston (2002).

³ Evans 1982: 112

⁴ Fodor 2001: 14

⁵ See chapter two for a more formal version of this argument.

⁶ Carston 2002: 58. Lepore and Loewer (1981) argue that translational semantics cannot do everything that a truth-conditional semantics can. For Carston, it is thoughts that provide this truth-conditional semantics.

sensitive.⁷ Natural language sentences are context-sensitive insofar as they do not determine which of a range of thoughts they are being used to express. On this view, taking the context into account to determine the proposition expressed by a representation is a matter of determining which thought it expresses.

This view of context-sensitivity entails that if a thought makes a reference to the context to determine its content then the only way in which to take the context into account is to generate another thought. If the context-sensitive thought was taken to be foundational then there cannot be a second thought that takes its context-sensitivity into account. If there was then the thought would not be foundational. So, either there cannot be any thoughts that are both foundational and context-sensitive or there are other means of taking the context-sensitivity of thought into account. If so, foundational thoughts can be context-sensitive. It is the second option that I will be considering.

This second option has some interesting implications. For one, thoughts are not necessarily the foundation of representation as on the MV. Instead it is thoughts and some other feature that form the basis of representation. What this other feature is varies with the view being considered but it will be some feature that refers to the context. Whether representations can function like this will be a central theme of this thesis. I will examine three different versions of context-sensitivity at the level of thought.

The first of these is the claim that there can be indexicals at the level of thought. In natural languages, such as English, these are expressions such “I”, “here” and “now” that vary in predictable ways with the context. For instance, who the speaker that uses “I” is will change the referent of “I”. In examining this I will focus on *de se attitudes*, which are attitudes that one has about oneself with the awareness that they are about oneself. That will be the topic of chapter 3. The second version of contextualism will be Searle’s view of intentionality, according to which we need many assumptions about the world to make a given thought express a proposition, which is covered in chapter 4. The

⁷ See Fodor 2001, Carston 2002

third is a judgement-dependent view according to which a concept needs a judgement to apply it to the world. These judgements are sensitive to contextual factors that I will elaborate on. I will consider this version of context-sensitivity in chapter five.

Another key concern will be whether any of these accounts of fundamental context-sensitive thoughts will be compatible with the computational theory of mind (CTM). I will elaborate on the CTM in chapter 2 but, for the moment, it can be summarised as the view that some cognitive processes that humans perform can be understood as computations. Compatibility with the CTM is important as it is the clearest and most promising attempt to understand our cognitive processes. I will argue that there is an understanding of the CTM on which the context-sensitivity of thought is possible in chapter 3.

The purpose of this chapter, however, is to introduce the notion of context-sensitivity in natural language (such as English, German and Japanese) and to state which understandings of context-sensitivity I will be investigating in this thesis. I will also relate these to some of the existing views on context-sensitivity in natural language. In particular, what I will call (following Hansen) Really Radical Contextualism (RRC).⁸ Contextualist views are those on which context-sensitivity extends beyond obvious cases such “I”, “here” and “now” to a broader range of representations. On RRC almost all natural language sentences and thoughts are context-sensitive. I will begin with some examples of context-sensitivity in a natural language (English) in S1.1. In S1.1.2 I will define some key terms that will be of use throughout the thesis and use these to explain the kinds of context-sensitivity that I will be focusing on. In S1.2 I will give an account of the degrees of context-sensitivity by briefly giving an account of the views of natural language which transition from a high degree of context-sensitivity in natural language, as in RRC through to low degree in Semantic Minimalism. Following that I will give an outline of the thesis that I will be defending in S1.4. I will also briefly explain why this thesis is significant for the current literature.

⁸ Hansen (unpublished): 2

1.1 What is Context-Sensitivity?

In this section I will give some examples of context-sensitivity in natural language to clarify the phenomenon. I will then introduce some terms that are useful in understanding the various accounts of context-sensitivity. Next, I will explain the types of context-sensitivity that I will be focusing on in this thesis. These are kinds of underdeterminacy that I will elaborate on below. They also give a sense of how pervasive the phenomenon of context-sensitivity might be, with views ranging from context-sensitivity as a pervasive phenomenon that permeates nearly all sentences to context-sensitivity as a less pervasive phenomenon which is present in only a few sentences. This will provide a useful point from which to begin discussing some of the views that are found in the literature.

1.1.1 Examples of Context-Sensitivity

A first example of context-sensitivity can be found in the sentence “I am happy”. This will express different propositions depending on who is saying it. If Sally says it then it would express the proposition *Sally is happy*, if John were to say it then it would express the proposition *John is happy*. So, it seems that the meaning (understood broadly, I will be more precise below) will vary depending on the context, in this case the person who is speaking.

Another example of the context-sensitivity is seen in the following example:

[Shortly after B first met Sam]

A: What did you think of Sam?

B: Sam is a saint.

A: No, Sam is not a saint, he is often rather unfriendly.⁹

Here it looks as though “saint” has not been taken literally, to mean an individual who has been canonized after dying. This is close to the literal meaning of “saint”. On the translational semantic account (see section 1) the literal meaning of the word type “saint” would be the concept type SAINT

⁹ Allott and Textor 2012: 187

(I will use caps to indicate a concept rather than a word). Instead “saint” here seems to communicate a different concept, like “Someone who is very kind and friendly”, which we might call SAINT*. Here it appears the context is playing a role and changing the meaning of “saint” from SAINT to SAINT*, loosely put (For the moment I am using “meaning” in an intentionally broad way. I will give a more precise account of terms below).

A third example is one in which two people are having a conversation as follows:

A: I am out of petrol.

B: There is a garage around the corner.¹⁰

Here, B seems to be telling A that she can get some petrol around the corner. However, it does not seem as though there is any part of B’s sentence which mentions this directly and if B had said it in different circumstances her utterance may not carry this meaning at all. (Imagine, for instance, that B had been asked to describe the surrounding area to someone who had just moved there).

These examples show that there are different sorts of context-sensitivity at work in language, and that they can be more or less closely related to the sentence itself. I take it that the first two are more closely connected to the sentence than the third example, as it is parts of the sentence that are being contextually altered or referring to the context in these cases. Whereas in the third case it seems that the words in the sentence could mean the same thing whilst the context-sensitive meaning picked up on here could vary. This shows that context-sensitivity is not limited to a single sort of phenomenon, and that there is a need to define some terms more carefully, most crucially “meaning”, to avoid confusion.

Broadly speaking it is possible to distinguish two kinds of meaning, one is the lexical or conventional meaning of a term and the other is the truth-conditional content of the term.¹¹ Returning to the

¹⁰ Grice 2010/1975: 177

¹¹ See Searle 1980, Travis 1997: 87, and Carston 2002 who all make this distinction in some form.

second example above, the lexical or conventional meaning of “saint” is someone recognised by the Christian faith as holy or virtuous. So, it might be conventional to use the word “saint” to express the concept SAINT. But in the context in which it was uttered the truth-conditional content was different, so that the utterance of “Sam is a saint” was true only if Sam was a SAINT*.¹²

In the third example given we seem to have a different sort of context-sensitivity at work. This form of context-sensitivity relates to what are typically called (following Grice) “implicatures”.¹³

Implicatures are propositions that are communicated indirectly by an utterance and do not involve a change to the content literally expressed by the utterance. In the next subsection I will give more detail on implicatures and the story that we might tell about this reasoning.

1.1.1.1 Implicatures

The exact definitions of implicatures vary but there are some important features which seem to remain constant. For instance, it seems that implicatures must always be calculable.¹⁴ That is, the hearer of the utterance must be able to work out what the implicature is. Furthermore, we, as theorists, should be able to provide a story about how hearers do calculate these implicatures.

Another feature that tends to remain constant is the cancellability of implicatures. So, B could deny that A could get petrol at the garage without explicitly contradicting herself. For instance, the utterance might continue “There is a garage around the corner, but I don’t think it is open now”. This cancels the implicature but nothing the speaker has said means she has contradicted herself. I will now give some detail on the reasoning that Grice thinks results in implicatures.

Grice’s reconstruction of the reasoning that leads to implicatures is based on the co-operative principle. The co-operative principle is ‘Make your conversational contribution such as is required at

¹² Recently, the definition of “saint” has been expanded to include this informal understanding of “saint” as SAINT*. This makes this example look more like a case of ambiguity than I would like as an illustration of context-sensitivity. Other examples of context-sensitivity include “This steak is raw” which is used to communicate that the steak is undercooked rather than raw.

¹³ See Grice 2010/1975.

¹⁴ Grice 2010/1975: 176

the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged'.¹⁵ For Grice, the utterances that trigger implicatures seem to fall short of this standard. Suppose an utterance with the content p is such an utterance. If the hearer thinks that the speaker is being co-operative, then they must have meant to communicate something else by their utterance. The speaker could reasonably expect the hearer to work out that they meant q based on p and q would be compatible with them being co-operative. Therefore, the speaker implicated that q by their utterance.¹⁶

1.1.1.2 Kinds of Implicature

There are three kinds of implicature that Grice discusses. One category is conversational implicature. A conversational implicature is generated by the utterance of a sentence in a context and the context is necessary for this implicature to be generated. There are two kinds of conversational implicature. The first is particularised conversational implicatures which are generated by making an utterance of P and 'in virtue of special features of the context, cases in which there is no room for the idea that an implicature of this sort is *normally* carried by saying that p '.¹⁷ In these cases the implicature depends on a feature of the context for its generation. The garage example is a typical example of that, and we can change the implicature by changing the context. So, if A had asked where they can buy a newspaper then B's utterance would implicate that A can get a newspaper at the garage. Arguably, no implicature would be generated if A had asked whether there was a garage around the corner, except perhaps that the garage was open. In these cases, special circumstances are needed for that implicature to be generated.

Second are generalised conversational implicatures in which the 'use of a certain form of words in an utterance would normally (in the *absence* of special circumstances) carry such-and-such an implicature or type of implicature'.¹⁸ One example is "X went into a house yesterday and found a

¹⁵ Ibid 173

¹⁶ Ibid 176

¹⁷ Ibid 180

¹⁸ Ibid 180

tortoise inside the front door”. Often, this would implicate that X had gone into a house that was not their own or the speaker’s.¹⁹ There is some dispute as to whether particularised and generalised implicatures constitute two distinct categories here.²⁰ Given that the difference between the two forms of conversational implicature is between the number of contexts in which it can be generated, the difference may be one of degree rather than kind.

Finally, there are conventional implicatures. In these cases, ‘the words used will determine what is implicated, besides helping to determine what is said’.²¹ One example of this is “Mary is poor, but she’s honest”. Here the “but” implicates that there is some tension between Mary’s being both poor and honest. However, it seems that a speaker can (just about) cancel this implicature without contradiction. Therefore, it is not a part of the meaning or content of that sentence.

Before giving a characterisation of context-sensitivity, it will first be necessary to introduce and define some terms. I will do this mostly via stipulation and none of the definitions that I will give will be without controversy. However, in most cases it should be possible to replace my preferred definition with one’s own preferred account without causing too much damage to the rest of the thesis.

1.1.2 Defining Terms

PROPOSITION – Propositions are the bearers of truth or falsity. It is the proposition that a given utterance expresses that is true or false. This can be understood as a function from possible worlds (and other indices) to a truth value, of either 1 (=true) or 0 (=false). ‘Thus an appropriate intension for a sentence is any function from indices to truth-values; an appropriate intension for a name is any function from indices to things; an appropriate intension for a common noun is any function from

¹⁹ Ibid 180

²⁰ Carston 2002: 111

²¹ Grice 2010/1975: 173

indices to sets.²² A proposition is a function from indices, here understood as a possible world and a time, to truth values. I will identify propositions with Kaplanian contents (see below).

There are other possible understandings of proposition that one could opt for. For instance, it could be understood as a Russellian proposition, i.e. as ordered pairs of objects and properties. However, I will focus on the functional understanding of propositions. This is because it can be used to provide a clear framework to determine where different sources of context-sensitivity might reside, as I will discuss in section 1.1.3.

SENTENCE IN NATURAL LANGUAGE – A sentence in natural language is a well-formed formula using the words of that natural language. A sentence is well-formed if it complies with the rules of grammar that govern that language.

CHARACTER – I will use Kaplan’s definition of character, which is stated as follows ‘...it is convenient to represent characters by functions from possible contexts to contents.’²³ Here the context is the context of use, the context in which the term is uttered. So, the word “I” has a character which often takes one from contexts of use to a content that is about the speaker. Indexicals such as “I” are examples of a context-sensitive character.²⁴ The content that they express is dependent in predictable ways on the context of use. For instance, “I” will reliably pick out the speaker when it is used.

Other words, such as “saint” will also have a character. These words will differ from indexicals and demonstratives in that they have a ‘fixed’ character.²⁵ In chapter four I will examine some ways in which these fixed characters might also admit of context-sensitivity. Nevertheless, I do not think that this would show these words to be indexicals. This is because the interaction with the context of use is much more complex than in the case of indexicals and demonstratives. The reason being that the

²² Lewis 1970: 23

²³ Kaplan 1989: 505

²⁴ Ibid 506

²⁵ Ibid 506

mechanisms by which these characters get their content will be very different to those of indexicals and they need not pick out any referent that is salient in that context. See chapter four for more discussion.

CONTENT – A content can be represented ‘...by a function from circumstances of evaluation to an appropriate extension. Carnap called such functions *intensions*’.²⁶ Here the circumstances of evaluation can be understood as possible worlds, much like Lewis’s indices, plus a time and place.²⁷ So we take an input (possible worlds) to a function (content) and give an extension (a truth-value). For example, if one entertains the thought that *there is a pizza in the fridge*, this will be true only of those possible worlds in which there is a pizza in the fridge. A circumstance of evaluation is the circumstance against which we assess the content (proposition), which may not be the same as the context in which the sentence was uttered (a context of use).

When it comes to thoughts the context of use will be the circumstances in which the thought is tokened. I will focus on occurrent thoughts as these have a clearer context of use than dispositional thoughts, though any given tokening of a dispositional thought might still be treated as an occurrent thought. Dispositional thoughts are a thought type that we are disposed to have, such as beliefs and desires.²⁸ When I speak about beliefs it will be as a token of a type of thought that we are disposed to have, but the focus will be on that token understood as an occurrent thought.

EXTENSION – The extension of a proposition is a truth-value, either true or false, for a given world, often the actual one. Possible worlds are paired with truth-values and thereby have an extension, for that content. The extension of a predicate is all the things that fall under that term.

²⁶ Ibid 502

²⁷ Lewis 1970: 24

²⁸ Smith 2005: 145

TRUTH CONDITIONS - For the purposes of this thesis I will understand truth conditions as follows. The conditions under which a proposition would give the value true. That is, the way the possible world must be for it to give the value true as an output for that function.

LEXICAL MEANING – Lexical meaning is a convention on the use of words, a guide to these can be found in a dictionary. The basis of this kind of meaning will be a convention concerning how those words are to be used in a given language L. The character of a word is a good candidate for the conventional meaning of that word. The lexical meanings of a word stay constant across contexts in their contribution to the meaning of a sentence.

A possible counter example to this claim are cases of lexical ambiguity in which a word, such as “bank” has two possible characters. In this case, it can be either a river bank or a financial institution. This appears to be a case of one word with two different lexical meanings. Alternatively, one can understand ambiguity as two distinct words that happen to be either homophones or be co-spelled, in which case ambiguity is not a counter example. In any case, it does not seem to damage this view if a word can have more than one character assigned to it on an occasion.

What meaning amounts to varies with the contextualist position. One common way of understanding meaning is as a logical form, which is often taken to be no more than a propositional form. Logical forms are defined as being the real structure of a sentence.²⁹ It is also ‘a well-formed formula, a structured set of constituents, which undergoes formal logical operations determined by its structure’.³⁰ According to some contextualists, a propositional form falls short of determining a proposition. The reasoning for this varies with the contextualist, but most who do accept that linguistic meaning alone is generally insufficient for a proposition. One strategy is to deny the existence of eternal sentences (sentences that have the same truth-conditions across all conditions of use).³¹ I will discuss others in chapters four and five. So, anything that is derived solely from the

²⁹ Stanley 2000: 392

³⁰ Sperber and Wilson 1995: 72

³¹ Carston 2002: 30 – 42

linguistic meaning of a term cannot be propositional (sometimes exceptions are made for analytic and mathematical sentences, but not always, see chapters 4 and 5). Instead it provides a frame from which to develop a proposition.

On non-contextualist accounts, the lexical meaning can establish a content with little contribution from the context. This is a more substantial, stronger, understanding of lexical meaning of words. I will make use of the weaker understanding of lexical meaning, according to which only a propositional form is determined, as it forms a common assumption between proponents of the MV and RRC. On its weaker understandings, a logical form may only determine a range of possible contents that are compatible with that sentence's propositional form.

THOUGHTS - The notion of thought that I will be using for this thesis will be Fodor's Language of Thought, sometimes called "Mentalese". I will elaborate on this more fully in chapter 2, but for the moment I will give a short account of thoughts. Thoughts are sentences in Mentalese.³² A sentence in Mentalese is composed out of concepts e.g. THE CAT IS ON THE MAT. These concept types can be physically realised in the brain as differing tokens. This follows from a token physicalist view of the mind, which I will elaborate on in chapter 2. Concepts are subject to computations in virtue of their logical form, and this allows for truth preserving inferences which preserve the semantic features of that sentence. These representations explain the psychology and behaviour of the individuals that entertain them. It is in virtue of these mental representations that we can entertain propositions.

These different notions fit together in the following way. A character, alongside a context of use, can determine content. A context of use is the context in which a sentence is uttered. This content is a function from circumstances of evaluation to extensions. A context of evaluation is the context under which we would assess the content for truth. This accounts for some forms of context-sensitivity in natural language. For example, the character of indexicals, such as "I", takes certain features of the context of use, in this case the person speaking, to give a content. There may also be instances of

³² Fodor 1975

context-sensitivity which are not due to the character of an expression. (This is one of the claims at the centre of the dispute between Minimalists and Contextualists). The change in meaning of “saint” from SAINT to SAINT* is not one that is mandated by the character of that expression. In the next section I will elaborate on some of the ways in which a proposition might be underdetermined by a representation.

1.1.3 Underdeterminacy

According to proponents of context-sensitivity either the content does not determine an extension, or the character of a term does not determine a content, without some input from the context. This results in three kinds of underdeterminacy that could be advocated by contextualists. Here I will introduce these terms and then discuss some potential problems with them, including whether they are all properly distinct from each other.

The first is character-underdeterminacy.

Character-Underdeterminacy: A structured representational item S is character-underdetermined if and only if a component of S makes an explicit reference to the context to establish content.

Examples include indexicals such as “I”, “here” and “now”. Here, the character requires input from the context to give any content as an output. The content given will vary with the context of use in predictable ways. For instance, the content expressed by “I” will vary with the speaker (when “I” is being used for direct reference). There is universal agreement that this kind of underdeterminacy is present in natural languages to at least some degree.³³

The second is:

³³ There are some who think that this kind of underdeterminacy explains many other cases of context-sensitivity in natural language. This view is indexicalism, which can be understood as the claim that ‘all truth-conditional effects of extra-linguistic context can be traced to logical form’ (Stanley 2000: 391). Here the logical form is intended to be real structure of that sentence, as something distinct from its surface grammatical form. I will not be considering the extent of this kind of context-sensitivity in thought. Rather I will consider whether this is possible at all at the level of thought, see chapter 3.

Type-Underdeterminacy: A non-indexical structured representational item *S* is type-underdetermined if and only if there are tokens of *S* that have distinct truth-values.³⁴

On this kind of underdeterminacy a sentence can be used to express at least two contents that do not have the same extension across all possible worlds, where this change in content is not due to character-underdeterminacy. Those who argue for this sort of position would likely claim that “Sam is a saint” is type-underdetermined as it can express at least two contents, that *Sam is a saint* or *Sam is a saint**.³⁵

The third is:

Token-Underdeterminacy: A token of a structured representational item *S* is token-underdetermined if and only if for some possible states of affairs its truth-value is indeterminate (i.e. if and only if it determines a partial function from possible worlds to truth-values).³⁶

In cases of token-underdeterminacy *S* might express only one content but that content will not determine a truth value for all possible worlds. A token-underdetermined representation will not be true or false for some possible worlds, though it may have a truth-value for others. For those cases in which *S* is token-underdetermined, determining a truth value may be a process that, if possible, depends on factors of the context. I will present such a view in chapter 5. Token-underdeterminacy differs from type-underdeterminacy because the content expressed by *S* can remain fixed, meaning

³⁴ Jaque 2017: 3. This has also been described as context-sensitivity by MacFarlane (2009: 232)

³⁵ It is sometimes not entirely clear whether the context-sensitivity of a term should be considered an instance of character or type-underdeterminacy e.g. the sentence “Fido is sitting”. This is because it may be unclear whether a given set of terms count as indexical or not. For instance, tense can be interpreted as either an operator on a sentence or as a variable that functions like other pronouns. Partee (1973) argues for the latter interpretation, which seems to make tense an instance of character-underdeterminacy. King (2003) also argues against a tense operator account. As such, I will tentatively treat tense as an instance of character-underdeterminacy, though I do not think that too much will hang on this.

³⁶ Jaque 2017: 17-18. MacFarlane (2009: 233) has described this as context-sensitivity. As has Recanati (1997: 90) who claims that ‘...a representation can be semantically indeterminate in two different ways: either it has a determinate character but that character fails to determine a content (in the context at hand); or it lacks a determinate character in the first place’.

that it is not type-underdetermined, whilst it may still be unclear what truth-value the proposition expressed by S should have.

An example of token-underdeterminacy might be the sentence “that is a cat” where the object in question appears to be a feline but also speaks Latin.³⁷ In this case it is unclear whether the object in question is a cat. This can be understood in an epistemological or a metaphysical sense. On the epistemological sense, there really is a determinate truth-value for that possible world but we are unable to work it out. On the metaphysical understanding, there is no such truth-value to be had in virtue of the world and the concept alone. I will use token-underdeterminacy in the metaphysical sense.

A fourth understanding of underdeterminacy is that ‘the literal meaning of an utterance underdetermines the speaker meaning’ where speaker meaning is something like the proposition(s) that the speaker intended to convey.³⁸ I will not be concerned with this claim. When it comes to (foundational) thoughts, there is an issue as to whether they could always express a proposition without a context. In respect of this issue, character, type and token-underdeterminacy seem to be more appropriate kinds of underdeterminacy than speaker-underdeterminacy. It is harder to understand speaker-underdeterminacy at the level of thought. Whilst there may be cases in which we may wonder why we have had a thought, it does not seem to be correct to attribute this to another speaker or to wonder what the thought intended to express as it is not obvious that there should be any communicative presumption made.

More significantly, speaker-underdeterminacy is an epistemic problem. It is a problem of trying to decide what the speaker intended to convey by their utterance given that what they have said is not isomorphic with the content they want to communicate. For example, the sentence “Billy is smart” might be used to communicate a range of propositions such as *Billy is smart enough*, *Billy is not*

³⁷ This is a case of Open Texture, which I will discuss in more detail in Chapter 5. See Margalit 1979 and Recanati 2004: 141 - 144

³⁸ Scott-Phillips 2015: 2

smart (in the case of irony), *Billy is smart for a ten-year-old* and so on. The sentence could have a determinate content and there would still be an additional concern about what the sentence was being used to convey by the speaker. It is not a metaphysical worry about the nature of thought. The central concern will be the other three kinds of underdeterminacy.

To sum up, an expression *e* is context-sensitive IFF at least one of the following conditions applies (i) *e* is character-underdetermined, (ii) *e* is type-underdetermined, or (iii) *e* is token-underdetermined or (iv) *e* is speaker-underdetermined and if at least one of (i) – (iv) is true there is a way in which the context can play a role in removing that underdeterminacy. If there is no way of removing that underdeterminacy then there is no context-sensitivity, only underdeterminacy. Condition (iv) is very rarely denied by anyone, even those who reject most forms of context-sensitivity do not deny that sentence meaning will typically be speaker-underdetermined.³⁹ So, it is the first three kinds of context-sensitivity which are of the most interest. It is also easier to understand how these might exist in thought.

1.1.3.1 Potential Issues with Underdeterminacy

Now that these versions of underdeterminacy have been introduced, I will consider some worries with them. One is that type and token-underdeterminacy fail to be distinct from each other. To see this worry, consider the role that the context plays in each case. In the type-underdeterminacy case, the context would act to decide which function, and therefore which extension, is the appropriate one. In the token-underdeterminacy case the context's role is to help fill out the partial function, resulting in a complete function or a value for that input. There may, after all, be several total functions that are compatible with a given partial function. Now, it appears, in each case, that we use the context to pick a function. In which case, we just have the same phenomenon in different guises. So, in cases of context-sensitivity the context contributes to picking a function, with the explanation for why this is the case being that we have either two or more possible functions to choose from, or

³⁹ Borg 2004: 255

we have only a partial function which is compatible with at least two other functions. In both cases, all that seems to occur is picking an appropriate function for the context.

In response to this, it is worth emphasizing that whilst both cases are similar in that the extension will be underdetermined in each case, the two still come apart in other ways. The sorts of phenomena that would count as a proof of each of one will not always count as a proof for the other. Whilst type-underdeterminacy can be understood as a sort of ambiguity (does this representation stand for this function or that function?) cases of token-underdeterminacy need not be ambiguous. So, we can know exactly which partial function it is that we are considering, and not know whether it is true in a given case. These are sometimes called cases of open texture.⁴⁰ As in the case in which a creature is indistinguishable from a domestic cat, except that it speaks Latin, or can be revived from death.⁴¹ We know that we are wondering whether that animal is a cat or not. We are not wondering whether we should apply CAT1 or CAT2. We are wondering whether CAT includes that animal or not.

Also, unpacking the role of the context in determining an extension will vary depending on which kind of underdeterminacy is in play. Whilst token-underdeterminacy can be resolved by selecting a correct compatible function, this is not the only way to do so. Instead, it may be that one can only have partial functions, and then we take a possible world as input which does not yield a truth value, we must decide how the partial function should be extended. Whilst this may determine the partial function for that input, it does not guarantee that the partial function has become a function. I will discuss this in chapter 5. So, the two kinds of underdeterminacy are distinct.

1.2 Views of Context-Sensitivity in Natural Language

Here, I will specify the degrees of context-sensitivity. This concerns how much of the language is context-sensitive. I will go from the highest degree down to the lowest degree. The view on which natural language is pervasively context-sensitive, as well as thoughts, is called Really Radical

⁴⁰ Margalit 1979

⁴¹ Recanati 2004: 142, Margalit 1979: 142

Contextualism (RRC). There are then those who maintain that only natural language is pervasively context-sensitive, called Radical Contextualists (RC). After these are Moderate Contextualists (MCs), who hold that there are some clear cases of context-sensitivity in the form of overtly context-sensitive expressions, as well as some covertly context-sensitive expressions. Finally, there are Semantic Minimalists (SMs) who maintain that context-sensitivity in natural language is exhausted by cases of overt context-sensitivity e.g. indexicals and demonstratives.

1.2.1 Really Radical Contextualism

Versions of RRC present in the literature include Travis's occasion sensitivity, as Travis claims 'Any representational form underdetermines when what has, or had, it would be true'.⁴² Another instance of RRC is put forwards by Searle, who claims that '...the features we have cited [dependence of content on the Background] are features not just of semantic contents but of representations generally, in particular they are features of intentional states, and since meaning is always a derived form of intentionality, contextual dependency is ineliminable'.⁴³ These are clear statements of RRC.

RRC breaks down into two key claims:

- a. All (or almost all) natural languages sentences are necessarily context-sensitive.
- b. All (or almost all) Sentences in thought are necessarily context-sensitive.

The possibility of claim (b) will be a major concern in chapters four and five. But for the most part I will be concerned with a weaker claim, that it is possible that some foundational thoughts could be context-sensitive.

From RRC we get the greatest degree of context-sensitivity available in the literature. On this view, all natural language sentences and thoughts that aim to express a proposition are context-sensitive.

An exception is made by Travis for mathematical expressions (and possibly some other abstract

⁴² Travis 2006: 32

⁴³ Searle 1980: 231. Whilst I will refer to this the Background Theory it is worth emphasizing that it is not a view in the background of Searle's other views. It is a position he explicitly argues for.

ones), however Searle claims that even these will be context-sensitive.⁴⁴ Given how pervasive this view takes context-sensitivity to be, the expression of a proposition is determinate only when there is a context. On these views a context is necessary for a representation to express a truth-evaluable content, though it need not be sufficient. I will now elaborate briefly on each version of RRC, with detailed presentations to follow in their respective chapters.

1.2.1.1 The Background Theory

Searle's view is that '...the meaning of a sentence only has application (it only, for example, determines a set of truth conditions) against a background of assumptions and practices that are not representable as a part of the meaning'.⁴⁵ According to Searle, this view is true of all intentional states. So, Searle adheres to both (a) and (b) of RRC. If Searle's view is even possibly correct this would be a point in favour of the view that all (or almost all) sentences in thought could be context-sensitive. If Searle's view is coherent then it shows how to make sense of context-sensitive thoughts. I will understand Searle's position as an instance of type-underdeterminacy.

The Background itself is a set of assumptions that are manifest to a speaker at a given time. For an assumption to be manifest it must be the case that the thinker can represent it and accept it as probably true. This makes the Background potentially very large. It also entwines our ability to express determinate content with our ability to assume things about the world. I will discuss this view in more detail in chapter 4.

1.2.1.2 Occasion Sensitivity

The second account of RRC that I'll consider is a judgement dependent view, largely following Travis. Travis claims that 'Any representational form underdetermines when what has, or had, it would be true'.⁴⁶ So, I will interpret Travis as an advocate of token-underdeterminacy, according to which the extension of a representation is determined by acts of judgement. For these judgements to be

⁴⁴ Travis 1989: 28, Searle 1980: 230

⁴⁵ Searle 1980: 221

⁴⁶ Travis 2006: 32

correct they must be in line with the judgements of a reasonable judge. So, Travis claims that ‘...for an item to have a semantic property P is for it to be so that a reasonable (informed) judge would take it to have P’.⁴⁷ These judgements can vary according to what the context demands and makes reasonable. I will discuss the judgement dependent view in chapter 5.

1.2.2 Radical Contextualism

Next there are Radical Contextualists (RC). RCs claim that (almost) all natural language sentences are context-sensitive. In holding this view they’re often taken to be committed to the following claims:

‘1 Linguistically encoded meaning underdetermines the proposition expressed by an utterance (its truth-conditional content).

2 Linguistic underdeterminacy is an essential feature for natural languages because there are no eternal sentences in natural languages.⁴⁸

Sperber and Wilson also maintain that ‘Contextual information is needed to resolve what should be seen as the semantic incompleteness, rather than the ambiguity, of the genitive [sentence uttered]’.⁴⁹ This goes with their conviction that natural language sentences encode a logical form, which needs enriching.⁵⁰ Recanati is also a proponent of this view, so this is not a view that is restricted to Relevance Theorists. Recanati claims that contextualism is a live option, according to which ‘it is speech acts, not sentences, which have a determinate content and are truth-evaluable: sentences themselves express a determinate content only in the context of a speech act’.⁵¹ This shows some of the chief proponents of RC and how they understand their view.

⁴⁷ Travis 1989: 48

⁴⁸ Carston 2002: 83

⁴⁹ Sperber and Wilson 1995: 188

⁵⁰ Ibid 182

⁵¹ Recanati 2004: 154

These quotes reveal the core of RC, as summarised in (a). Where there is a debate as to whether a given kind of context-sensitivity is a case of speaker-underdeterminacy or a case of character, type or token-underdeterminacy, RCs tend to argue that it is one of the latter three categories.

1.2.3 Moderate Contextualism

Moderate Contextualists (MCs) claim that there is more context-sensitivity than is normally thought in natural language (i.e. more than overtly context-sensitive expressions such as indexicals and demonstratives), but that it is not a thoroughly pervasive feature. This degree of context-sensitivity can be understood as including members of a Basic Set of expressions that are widely agreed to be context-sensitive as well as a few additional items but not the entirety of the language. Members of the Basic Set include indexicals and demonstratives.⁵² MCs deny that (a) is true whilst still accepting that context-sensitivity is a widespread feature. MC is summarised in the following three claims:

‘(MC1) The expressions of the Basic Set do not exhaust all the sources of semantic context-sensitivity.

(MC2) Many sentences ... fail to have truth conditions or to semantically express a proposition; they express only fragmentary propositions... it is not truth evaluable.

(MC3) For the cases in question, only their utterances semantically express a proposition, and have (interpretive) truth conditions, and so, take a truth value’.⁵³

One example of MC is that what counts as knowing that P will be subject to differing standards depending on the circumstances (whilst maintaining that a basic set of natural language expressions are context-sensitive). This is epistemic contextualism. This is the view that ‘A single knowledge ascribing sentence can bear different truth values relative to different contexts of utterance, where this difference is traceable to the occurrence of “know,” and concerns a distinctively epistemic factor’.⁵⁴ This is an MC view as they have included not only items that are overtly context-sensitive or

⁵³ Cappelen and Lepore 2005: 7

⁵⁴ Schaffer and Szabo 2013: 492-3

a part of the basic set, but also words such as “knows”. As “knows” is not overtly context-sensitive, this makes the view contextualist. However, they are not RCs because nothing in their view commits them to (a).⁵⁵ I will not discuss MC as such in this thesis.

1.2.4 Semantic Minimalism

Penultimately, the view that allows for the lowest degree of context-sensitivity in natural language is Semantic Minimalism (SM). There are two main versions of this view. One is from Borg, who maintains that ‘there are propositional, truth-evaluable contents which attach to all well-formed declarative sentences (relativized to a context of utterance), in virtue of the standard lexico-syntactic constituents of those sentences alone’.⁵⁶ The other comes from Cappelen and Lepore, who also maintain that sentences can express a proposition without assistance from the context in conjunction with what they call speech act pluralism (I will give an explanation of this below).⁵⁷ The differences between these versions of SM are interesting, but will not be too significant for the following discussion.

Here I will present some of the key features of both versions of SM. SM only allows for a very small set of context-sensitive expressions of natural language. It is just that where there is context-sensitivity, it is demanded by either the syntactic structure or the lexical content of the expression.

Borg’s version has four key claims:

- i) Semantic content for well-formed declarative sentences is truth evaluable content.
- ii) Semantic content for a sentence is fully determined by its syntactic structure and lexical content.
- iii) There are only a limited number of context-sensitive expressions in natural language.

⁵⁵ Whilst Cappelen and Lepore argue that MCs are committed to a RC view, we can for the moment suppose that they are not so committed for the purposes of exposition. Borg (2012: 29 – 43) also argues that there is not really a middle ground between RC and SM.

⁵⁶ Borg 2017: 1

⁵⁷ Though it is not clear that Cappelen and Lepore are both currently advocates of Semantic Minimalism, their view remains important.

iv) Recovery of semantic content is without access to speaker intentions.⁵⁸

(iv) is reconciled with cases of overt context-sensitivity such as “That is no good” by claiming that the task of discovering what the referent of “that” is a non-semantic matter. It is claimed that ‘...if we allow that that the semantic content of a demonstrative can be grasped without this entailing non-linguistic identification of the referent then we will be left with a fairly thin notion of semantic content here.’⁵⁹ The idea here seems to be that we can get some truth-conditions, and so a complete semantic content, even when we do not know (in any substantive sense) what the “that” in question is. It is enough that we know that the utterance would be true so long as whatever the object referred to is no good.

SM is also not concerned with giving an account successful communication; they concede that communication will require pragmatic processes.⁶⁰ However, they maintain that truth-conditional content can be had with only disambiguation and saturation. This SM view allows for character-underdeterminacy but only in a limited number of cases. It does not allow for type or token-underdeterminacy.

This view has recently been developed further so that the minimal semantic content of an uttered sentence is what the utterer can be held liable for.⁶¹ This is spelt out as ‘Judgements of a speaker’s liability or culpability for content (strict): a binary notion whereby a speaker A either is or is not held liable for a given content by their utterance of s’.⁶² This changes some of the focus on minimal content as a socio-linguistic concept, rather than just as a psychological or semantic one. This understanding of SM helps to clarify the notion of minimal content, as what one can be held

⁵⁸ Borg 2012: 4-5

⁵⁹ Ibid 138

⁶⁰ Borg 2004: 259

⁶¹ Borg 2017: 8

⁶² Ibid 8

responsible for in a legal context based on their utterance. However, it also makes SM less useful for this thesis as it is not clear that we can apply this kind of standard to thoughts.

The second version of SM is advocated by Cappelen and Lepore. For current purposes it is like Borg's version of SM except that context-sensitivity is limited to what they call the Basic Set and they endorse speech act pluralism. I will go over each in turn. The Basic Set includes 'personal pronouns... the demonstrative pronouns... the adverbs "here", "there", "now", "today", "yesterday", ..., "hence(forth)"... the adjectives "actual" and "present"... the contextuials, which include common nouns like "enemy", "foreigner", "immigrant"... as well as common adjectives like "foreign", "local", "domestic"...'.⁶³ The members of the Basic Set all seem to make an obvious appeal to context in order to communicate a proposition. For instance, if someone says, "He is going to that local post office now", I will need to appeal to the context of use to know who he was and when he was going. These all seem to be cases of context-sensitivity that are demanded by the lexico-syntactic features of the utterance, rather than from a feature of the context or from a general short coming of representations themselves.

Cappelen and Lepore combine SM with what they call speech act pluralism, the idea is that 'No one thing is said (asserted, or claimed, or ...) by any utterance: rather, indefinitely many propositions are said, asserted, claimed, stated, etc.'⁶⁴ Amongst these propositions will be the minimally enriched semantic claim. This will be one in which there need only be the resolution of ambiguities and the saturation of members of the basic set to get a proposition. Others will be more enriched and include instances in which "Saint" is modulated to assert that someone is very kind.

There may be a comparable view to SM at the level of thought according to which some of the components of thought are context-sensitive. These components may be equivalents of the Basic Set or overt indexicals. This will give us an idea as to whether there are indexicals or demonstratives

⁶³ Cappelen and Lepore 2005: 1

⁶⁴ Ibid 199

present in thoughts. For example, whether we could have a concept defined in an ostensive manner using demonstratives, such as *THAT SHADE OF RED* or the indexical thought that *SHE IS IN READING*.

The issue with such thoughts will not be to prove that we could have these thoughts. It seems possible that we do have such thoughts. Rather, the issue will be to try to determine whether any of these thoughts are foundational and context-sensitive. If we have these thoughts, then they may seem to require a process of saturation, where saturation is a syntactically mandated process for the provision of elements needed for the recovery of a proposition or a set of truth conditions.⁶⁵ So when it comes to *SHE IS IN READING*, it may be that the “*SHE*” needs to be replaced by some other content e.g. *KIM*. If these indexical, or demonstrative, elements can always be replaced, or are only true in virtue of thoughts which are saturated being true, then these will not be foundational thoughts. The thoughts that are foundational will not be context-sensitive.

A second issue is that computational operations are typically understood as operating on the properties of types, rather than on the properties of tokens. This poses a problem for indexicals and demonstratives at the level of thought, as it seems that they become useless when they cannot be treated as tokens.⁶⁶ Indexicals at the type level only have a character, but no reference. Only particular tokens of indexicals can have a truth value.

1.2.5 No Context-Sensitivity

There is, finally, a position that is totally unoccupied when considering natural language, which is that context-sensitivity does not exist. This theory is, when applied to natural language sentences, certainly false. Cases of indexicals and demonstratives seem to be too compelling a counter example for any such view to gain traction. However, when it comes to thoughts, it may be that there are no indexicals or demonstratives. It may be that each thing gets its own symbol and never varies in content or extension. If there are no forms of context-sensitivity at the level of thought, then

⁶⁵ Borg 2004: 46

⁶⁶ Ball 2010

thoughts will not be context-sensitive. There would be no sentences in thought that are context-sensitive. If this is necessarily correct, then the thesis will be false. Proponents of the MV, which I will discuss in more detail in chapter 2, maintain that natural language is context-sensitive whilst denying that there can be any context-sensitivity in thought. Proponents of the Mixed View will provide the main foil that I will use throughout this thesis as proponents of the Mixed View make the most prominent objections to the idea of context-sensitive thought, including the regress argument that I will discuss in chapter 2.

Proponents of the MV have been RCs. Where there is a choice between interpreting a case of context-sensitivity as a case of either speaker-underdeterminacy or a case of type, token or character-underdeterminacy, RCs tend to take the latter approach. As a result of this I will often assume that (a) is true as it is a point of common ground between proponents of the MV and RRCs. So, SM and MC will not be as relevant for this thesis.

1.3 Near-Side vs. Far-Side Pragmatics

The term “pragmatics” can be used to describe the ways in which the context is considered to remove underdeterminacy. This is needed to provide a means of removing the underdeterminacy that does appear, as in the second condition for a sentence’s being context-sensitive. There are two broad kinds of pragmatics in this sense. One is called near-side pragmatics; the other is far-side pragmatics. Broadly speaking, ‘Classical Gricean pragmatics, aimed at computing implicatures, is usually conceived as dealing with far-side pragmatics. It involves reasoning about why what was said, was said. Near-side pragmatics, on the other hand, is pragmatics in the service of determining, together with the semantical properties of the words used, what was said’.⁶⁷ So, of the examples given, “I am happy” is a case of near-side pragmatics. “I” needs to be saturated for there to be a complete content here. Here the pragmatic work is required to determine what was said by the speaker.

⁶⁷ Korta and Perry 2008: 349. Of course, for some of the instances of what I will include as examples of near-side pragmatics will also require reasoning about why a person has said what they did say.

By contrast, implicatures are cases of far-side pragmatics. To derive a conversational implicature one usually needs to reason about why someone said what they did. The same could be said of generalised conversational implicatures. In these cases, one might reason to the effect that the speaker would only say that if they wanted to express the content that is usually implicated by utterances of these sentences. The same applies to conventional implicatures. Whilst these are cases in which there is something in the sentence that triggers the implicature, one might suppose that they chose those words to generate the implicature.

I will focus on cases of near-side pragmatics at the level of thought. This means that my concern will be with what a given thought is used to think and not with why a given thought was tokened. This means my main concern will be with cases of thoughts that are like “I am happy” and “Sam is a saint” rather than cases like “There is a garage around the corner”. I will not consider the possibility of a far-side pragmatics of thought when asking whether there can be context-sensitive thoughts. This is partly because these sorts of cases are not considered in the literature to my knowledge and partly because cases of implicature seem to require communicative principles. It is not obvious how to apply these to thought and how we might distinguish implicated thoughts from those that merely follow in a chain of reasoning.

So, I will aim to give accounts of context-sensitivity of thought that exhibit either character, type or token-underdeterminacy and which make use of a near-side pragmatic process to make the proposition expressed determinate. This means that the processes involved will all make use of mechanisms that determine what was “said” by a thought rather than reasoning about why that thought was tokened.

1.3.1 More Examples of Context-Sensitivity

I will now unpack some examples of context-sensitivity, I will relate these to the kinds of underdeterminacy where possible. This should clarify the kinds of phenomenon that I am interested in and those that I am not. The examples that I will give are the following:

1. Multiple encodings (i.e. ambiguities)
2. Indexical references.
3. Missing constituents (or unarticulated constituents).
4. Unspecified scope of elements.
5. Underspecificity or weakness of encoded content.
6. Overspecificity or narrowness of encoded content.⁶⁸
7. Irony
8. Metaphor
9. Metonymy
10. Polysemy
11. Bridging inferences

I will elaborate these in turn.

(1) Ambiguity is when there are ‘...words with a single orthographic/phonetic type but with multiple meanings’.⁶⁹ A classic example of this is the English word “Bank” which can mean either a financial institution or the side of a river. A sentence of English that uses this word will therefore be type-underdetermined. Alternatively, ambiguity might be understood as there being two words, each with their own meaning.⁷⁰ In which case, a given representation is not type-underdetermined when it is ambiguous, as there are two distinct representation types in play. I will opt for the former reading as it is close to the reading RCs want to give. For instance, Carston calls ambiguity a case of ‘multiple encodings’.⁷¹

(2) Indexicals include words such as “I” “here” and “now”. ‘Indexicals have a *context-sensitive* character. It is characteristic of an indexical that its content varies with context’.⁷² These are words

⁶⁸ Carston 2002: 28

⁶⁹ Borg 2004: 141

⁷⁰ Sennet 2016: S3.1, Allott 2010: 146

⁷¹ Carston 2002: 28, Elbourne (2011: 34), whilst not a RC, also uses this account of ambiguity.

⁷² Kaplan 1989: 506

which have a character that ties their reference to the context in which they are made. These introduce character-underdeterminacy into the sentence in which they appear.

(3) An unarticulated constituent is, at first glance, a part of a sentence that appears in the proposition expressed but does not appear in the sentence. Examples include “It is raining [here]” and “Ibuprofen is better [than paracetamol]”. Unarticulated constituents (UCs) are understood in a few different ways in the literature.

One interpretation of UCs is that ‘...there are no truly unarticulated constituents, because, although some constituents do not appear in the surface syntax of the linguistic expression used, so are not perceptible, they are nonetheless articulated in the logical form of the expression’.⁷³ On this view, UCs are like indexicals. Stanley is an advocate of this view, called indexicalism.⁷⁴ On this understanding UCs are cases of character-underdeterminacy. This view strains the use of “indexicals”. Prototypical examples of indexicals vary in predictable ways with the context. For instance, “I” typically varies with the speaker and “you” with the addressee and “here” with the place of utterance. It is not clear that the other examples of UCs react to the context in this way, especially if this is meant to explain the change from, for instance, *saint* to *saint**.

On another interpretation a UC ‘...is an element which appears at the semantic level but which does not appear anywhere at the syntactic level – it is semantically relevant though syntactically unarticulated’.⁷⁵ On this view, UCs seem to be indistinguishable from free pragmatic effects. A free pragmatic effect is not mandated by any linguistic fact about a sentence but rather by considerations about how the language is used.⁷⁶ On this understanding of UCs there are two plausible kinds of underdeterminacy that might be applied here. One is speaker-underdeterminacy and the other is type-underdeterminacy. Speaker-underdeterminacy is plausible because it seems that these are

Borg 2004: 141

⁷⁴ Stanley 2000: 391

⁷⁵ Borg 2004: 213

⁷⁶ Recanati 2004: 10

cases in which a speaker is intending to communicate more than what they have said. This view is especially plausible if one thinks that the sentence as it stands can express a proposition with minimal assistance from the context. In this case, the UC is only needed to account for the speaker's intended meaning.

(4) Concerns the scope of elements in a sentence. One example of this is "Everyone isn't hungry", which is ambiguous between "Not everyone is hungry" ($\sim (\forall x) (Hx)$) and "No one is hungry" ($(\forall x) \sim (Hx)$) depending on the scope of the negation.⁷⁷ This is another instance of type-underdeterminacy as there is one sentence that can express at least two different contents.

(5) Underspecificity or weakness of encoded content. In these examples there is some element of the sentence which falls short of the intended meaning. This has been understood as the encoded meaning of a word covering too many items in its set. Alternatively, it could be understood as the intension being too coarse grained and therefore covering too much. An example of this is the sentence, "There is nothing on telly tonight".⁷⁸ Here, "nothing" is too broad as it seems that the intended content is that there is nothing worth watching on television. This is a free pragmatic effect as well, however it appears to be a case speaker-underdeterminacy.

(6) Overspecificity or strength of encoded content. In contrast to (5), this is when the set has too small an extension for the intended meaning, or the intension is too narrow. The altering of SAINT to SAINT* is an instance of a word's meaning being too specific for the content that the speaker intended to communicate. SAINT* might be an ad hoc concept, formed by a process of modulation.⁷⁹ This is also a free pragmatic effect that does not seem to be required by anything in the character of that sentence. This is another example of a controversial kind of case with some arguing that this is

⁷⁷ Carston 2002: 24

⁷⁸ Ibid 26

⁷⁹ See Wilson and Carston 2007, Allott and Textor (2012).

only a case of speaker-underdeterminacy whilst others take it to be a case of type-underdeterminacy.⁸⁰

(7) Irony can be understood as expressing a meaning that normally signifies the opposite of what the speaker intended to express. For example, suppose that Jane has just entered her office whilst it's raining heavily outside and says to their colleague "Lovely weather we are having". Of course, the weather is not good, and the speaker means to communicate this. This is speaker-underdeterminacy.

(8) The definition of a metaphor is a controversial topic. However, it is widely agreed that metaphors are not instances of literal meaning and instead have a lot to do with how a hearer interprets the utterance and what associations they have with the content. Metaphors also often apply predicates that are in some ways not literally applicable to the subject. For example, "Achilles is a lion" or "Juliet is the sun". It is often thought that these should be spelled out in terms of a far-side pragmatics, so metaphors are taken to be a special type of implicature.⁸¹ This is speaker- underdeterminacy.

(9) 'A figure of speech, metonymy is the use of a property to refer to its possessor, or an object associated with it'.⁸² One well known example of metonymy is "The ham sandwich wants her bill". In this case, the property of having purchased a ham sandwich is used to refer the person who has purchased the sandwich. This is speaker-underdeterminacy.

(10) 'When a word has more than one related sense it is said to be polysemous'.⁸³ An example of this is the word "book" which can have several related meanings. So, if I utter "That's my book" then I may mean that it is a book that I own, a book that I have written or a book that I am making use of.⁸⁴ In some of these cases there is also type/token distinction in play, so that if it were a book that I had

⁸⁰ This divide very often follows one's views on the extent to which one takes context-sensitivity to be prevalent in natural language. Those who think that Semantic Minimalism (see below) is plausible tend to argue that these cases are cases of speaker-underdeterminacy. Those who do not argue that these are cases of type-underdeterminacy, whilst indexicalists argue that they are cases of character-underdeterminacy.

⁸¹ Carston 2002: 378

⁸² Allott 2010: 125

⁸³ Ibid 146

⁸⁴ Pietroski 2018: 5

written then I am likely referring to the book type. If it is a book that I own then it may be either a token (so that I own that particular book) or a type, if I encounter another copy of a book that I already own. This is different from cases of ambiguity, as ‘... in ambiguity the different senses are unrelated. Lexical ambiguity may be described as several words that happen to be pronounced the same way; polysemy is one word with more than one sense...’.⁸⁵ Like ambiguity, this appears to be a case of type-underdeterminacy.

(11) Bridging inferences make use of general or common knowledge between interlocutors to make an inference about the intended meaning. They are often taken to be performed on sentences that express a content. For example, “Mary took out her keys and opened the door” is often taken to mean that Mary took out her keys and used them to open the door. Here the general knowledge is that keys are used to unlock or open doors. In this sense bridging inferences are either taken to be a free pragmatic effect as they ‘introduce syntactically unmarked information from the context of utterance’.⁸⁶ On this understanding they are cases of type-underdeterminacy. Alternatively, they may be a form of implicature, in which case they are more like cases of speaker-underdeterminacy.

A notable feature of this list is that none of these cases are obviously instances of token-underdeterminacy. Some RRCs argue that there are cases of representations expressing a partial function, independently of whether the expression has a character that explicitly references the context. The claim that representations express only partial functions is a radical claim. If it is true of the concept of, for instance, CAT then there could be cats which do not fall under that concept. This may be because the cat speaks Latin, so that we are not clear on whether CAT applies to it.⁸⁷ I will elaborate on this in chapter five.

⁸⁵ Allott 2010: 146

⁸⁶ Borg 2004: 47

⁸⁷ Recanati 2004: 142

1.4 Scope of the Thesis

Now that some of the degrees of context-sensitivity have been given, we are able to see what the scope of the thesis might be. There are stronger and weaker versions of the claim that there is no context-sensitivity in thought. The stronger version of this negative thesis is easy to formulate. I will begin by giving this antithesis giving some claims that run contrary to this claim and explain the versions that I will be most concerned with proving.

1.4.1 Antithesis

To begin with, the antithesis goes as follows:

Strong no context-sensitivity: There can be no sentences in thought that are context-sensitive.

Proponents of the MV maintain that this is so. There is a weaker version of this claim that is harder to understand but which is a more useful foil for the claim that I want to make.

Weak no context-sensitivity: There can be no context-sensitivity at the foundational level of thought.

Here, I intend to use “fundamental” (or foundational) in Evans’s sense, so that ‘... a level of thought which is fundamental in this strict sense: every thought about *G*s which is not of this level is conceived to be made true by the truth of thoughts which are of this level’.⁸⁸ Here “*G*” stands for any arbitrary property of being *G*.⁸⁹ A key component of a fundamental thought will be that the person having it will have a foundational idea of the object that is *G*. A foundational idea of an object is to have a ‘general conception of the way in which *G*s are distinguished from one another, and from all other things’.⁹⁰ I will not be committed to a strong account of what foundational ideas are meant to be like. It may not always be possible to distinguish whether an object is a member of *G* or not. It should be possible in cases lacking unusual features (e.g. poor lighting or a number of objects that are designed to fool people into thinking that they are members of *G* when they are not).

⁸⁸ Evans 1982: 112

⁸⁹ Ibid 104

1.4.1.1 Foundational Thoughts

This idea of a foundational thought should be clarified. Consider the thought “Pineapples and bananas are both fruits”. In this case, the truth of this thought depends on the truth of each of its conjuncts, e.g. that the thoughts that “Pineapples are fruits” and “Bananas are fruits” are both true. Does it follow that the conjuncts are fundamental to the thought that both bananas and pineapples are fruits? Initially, there does not seem to be anything in the definition of fundamental thoughts to rule this out. So, it seems that complex thoughts are dependent in various ways upon the simpler thoughts of which they are composed. These relations may well be like those that are familiar from logic. It may follow from this that negations are always foundational. For instance, $\sim p$ does not depend on the truth of p . Rather; p must be false for $\sim p$ to be true.

Another possible case is presupposition. For instance, the thought that “I intend to go to Paul’s party” might depend on it also being true that “I believe that Paul is having a party”. If I did not believe that this was true, then it seems that I cannot intend to go to the party. So, it seems to be a necessary condition for the thought that I INTEND TO GO TO PAUL’S PARTY that I also believe that Paul is having a party. Is the belief that I intend to go to Paul’s party less fundamental considering that it presupposes that there is a party? It seems not, there are some good reasons to reject this presuppositional account of what it is for a thought to be foundational.

For one, it seems that it is possible for this to trigger a regress for a great many kinds of thought, as any given thought may need to presuppose something to be true. Exceptions to this would be tautologies, which are always true. They do not depend on a presupposition to be correct. However, it does not seem accurate to say that only tautologies can be fundamental. It may become difficult to establish, in non-trivial cases, which presuppositions are the correct ones. There might be several possible sets of presuppositions for a given proposition, yet they may not all be compatible with one another. For example, “It is true that one should not commit murder” might presuppose that there are moral facts. However, there might be disagreement about how these facts should be understood. Whether moral facts should be understood in a deontological or a consequentialist manner will

affect which presuppositions we make about them. As these two views seem to be incompatible then it matters which we choose, and this selection will be important.

Another way to rule out presuppositions as fundamental thoughts is to say that fundamental thoughts should contain fundamental ideas. The worry with presuppositions is that they do not get one closer to a fundamental idea of the object. Taking the example of “I intend to go to Paul’s party” the presupposition that Paul is having a party does not get us closer to being able to identify that party. It is more appropriate to take the fundamental ground of difference to be an important part of what makes a thought fundamental, rather than a lack of presuppositions. This suggests that listing presuppositions will not take one closer to fundamental thoughts.

The interest in foundational thoughts, rather than thoughts considered more broadly, is important for several reasons. Firstly, context-sensitivity in thought would not be that interesting if it can be removed or replaced by non-context-sensitive thoughts. The thesis that thoughts are context-sensitive in a way that is resolvable is not so interesting when compared to cases of fundamental context-sensitive thoughts.

Secondly, proponents of context-sensitivity in thought often claim that ‘representational forms’ or ‘intentional states’ are unavoidably context-sensitive.⁹¹ So, it is closer to take these views as concerning foundational thoughts.

1.4.2 Thesis

This provides a few different versions of the positive thesis that I could defend. They go as follows:

Strong context-sensitivity (modal): All (or almost all) sentences in thought *could* be context-sensitive.

Strong foundational context-sensitivity (modal): All (or almost all) foundational sentences in thought *could* be context-sensitive.

Weak context-sensitivity (modal): Some sentences in thought *could* be context-sensitive.

⁹¹ Travis 2006: 32, Searle 1989: 231

Weak foundational context-sensitivity (modal): Some foundational sentences in thought *could* be context-sensitive.

Strong context-sensitivity (actual): all (or almost all) sentences in thought *are* context-sensitive. (This appears to be the view of RRC).

Weak context-sensitivity (actual): Some sentences in thought *are* context-sensitive.

Strong foundational context-sensitivity (actual): All (or almost all) foundational sentences in thought *are* context-sensitive.

Weak foundational context-sensitivity (actual): Some foundational sentences in thought *are* context-sensitive.

My main aim will be to at least prove that *Weak foundational context-sensitivity (modal)*, which is the claim that some foundational sentences in thought could be context-sensitive, is true. This can be done by showing that any of the three kinds of underdeterminacy (token, type and character) can be true of some foundational thoughts without contradiction. In chapters four and five I will be considering a stronger thesis, namely *Strong foundational context-sensitivity (modal)* which is the claim that all (or almost all) foundational sentences in thought *could* be context-sensitive. I aim to do enough to motivate further consideration of these views considering that they are not impossible, as has sometimes been claimed.

The upshot of this thesis comes from its disproving the MV and from its paving the way for a proof of the actual theses. Proving *Weak foundational context-sensitivity (modal)* also involves overcoming some of the objections to actual context-sensitivity, as most of these objections aim to conclude that such context-sensitivity is impossible.

This is in opposition to the MV.⁹² This is in keeping with a literature criticising the MV.⁹³ Where I believe this thesis differs from this literature is in its consideration of some hitherto rarely considered arguments, such as a regress argument that is taken to apply to context-sensitive thoughts, and in the consideration of ways in which thoughts could actually be context-sensitive.⁹⁴ So this attempts to show what a RRC view could look like whilst showing that it does not fall into the objections that proponents of the MV make.

1.4.2.1 Significance of Thesis

There are some other factors that seem to depend on whether this thesis is correct or not. These include things like whether thoughts follow a principle of composition, according to which the components of the thought, and the order in which they are assembled is all that is relevant to determining the content of that thought. There are, however, other understandings of composition that might be in play. For instance, it might be that the character composes but not the content. Or it may be that composition is relevant only to saturated elements, or elements that have already taken the context into account.⁹⁵ The most significant understanding of the composition principles here seems to be whether the content depends solely on the parts of the thought and the order in which they are assembled. If the context plays a necessary role in determining content, or extension, then it seems that the composition principle is threatened.

Also, at stake is the identity of concepts and whether thinking requires a connection to action. This also relates to issues as to whether a thought has its content essentially, or whether a thought requires a context to have a content or extension. If it is the latter, then a thought does not essentially express a function, but at best a partial function. This would be a blow to the MV. This is largely taken to be dependent on whether the truth-conditional principle of composition holds for

⁹² Jaque 2017: 3

⁹³ See, for instance, Jaque 2017, Elguardo 2005, Clapp 2013.

⁹⁴ An exception to this is Jaque (2017b) who acknowledges the regress argument and argues against it in favour of a RRC view.

⁹⁵ Pagin and Westerståhl 2011: 34

thoughts or not. These two are the most interesting and perhaps the most apparent consequences that this thesis might have.

1.5 Summary

In this thesis I will be concerned with showing that the *weak no context-sensitivity* view is false, and I will do so by attempting to show that there are some foundational thoughts that could be context-sensitive. The thesis that I will defend is that foundational, declarative sentences in thought (in a sense to be elaborated on in the next chapter) can be context-sensitive. I will do so by showing that at least one of the kinds of underdeterminacy are possible at the level of foundational thought.

1.5.1 Chapter Outlines

Chapter two will go into more detail on the topic of “thoughts”. Doing so will include giving an account of the CTM. I will draw on Fodor’s understanding of these topics. The CTM has some advantages as it gives a clear research project for cognitive science and helps to provide a way of individuating concepts and thereby thoughts. It will also include some discussion of the MV and lay out some of the arguments in favour of this view that have not yet been discussed by many advocates of non-Mixed Views. This concerns the regress argument against context-sensitive thoughts.

In presenting these views I will introduce several arguments that will be important throughout the thesis. Of major importance will be the Frame Problem, sometimes called the Relevance Problem, in computational accounts of reasoning. I will discuss and defend a version of the Frame Problem. I will do so to support an understanding of the CTM on which there are informationally encapsulated modules that operate computationally whilst there is a general reasoning processor that is not encapsulated, and which does not operate computationally. Motivating this general processor will be important as it provides a way of making versions of context-sensitive thoughts compatible with the CTM when they otherwise would not be.

In chapter 3 I will consider *de se* attitudes as a candidate for thoughts that are character-underdetermined. I will first follow Ninan (2016) in arguing that *de se* attitudes are essential to certain sorts of explanation of behaviour.

It can then be argued that indexicals are required for making sense of this essential feature of the context. In which case, it seems that we must have some degree of context-sensitivity in thought. But context-sensitivity in thought appears to be in tension with the computational theory of mind (CTM).

In this case we have an apparently inconsistent triad between the following three claims:

(i) De Se thoughts are essential.

(ii) De Se thoughts are indexical, they have a (Kaplanian) character.

(iii) Computations can only take the syntactic *type* into account, they cannot take tokens into account.⁹⁶

(i) and (ii) together seem to imply that we can have thoughts which are context-sensitive and that these thoughts cannot be replaced without a loss of content. The problem comes with claim (iii). If this is correct, then we cannot make sense of a thought which uses a character such that its referent could vary from tokening to tokening. This leaves us with a choice of which of these claims to reject. Rejecting claim (i) seems to be a move towards the inessential indexical view. Rejecting (ii) might leave us with a view akin to Millikan's (1990) on which *de se* thoughts are present but make use of non-indexical components. Rejecting either (i) or (ii) will be a move against the essential *de se*. Rejecting (iii), on the other hand, might require us either to reject the CTM or to explain how the CTM can be compatible with indexicals. This chapter argues for the second disjunct, showing how computations might be taken to be sensitive to character. This supports character-underdeterminacy in at least some cases of foundational thoughts.

In chapter 4 I will argue that it is unclear how the Background of assumptions is able to perform its role of allowing a sentence to have truth conditions. This is made particularly clear when considering

⁹⁶ Ball 2010

that adding new content that was in the Background, to provide a new thought, leads to a regress. That is, if one assumes that a thought cannot determine its content then it does not matter how much of the Background is added to provide a new thought, it will not provide determinate content. This is because that new thought will also depend on a Background of assumptions to determine a content. I will return to this regress argument in chapter two.

In response to this worry, I will argue that the Background is able to affect a concerning relation between a representation and its content. This means that the Background can play a role in determining content without having to add additional content, thereby avoiding a regress. It is enough that the Background can play the role of causing a concerning relation.

I will suggest using the concerning relation as a way of understanding Searle's Background theory, which provides an example of type-underdeterminacy being resolved that does not require the generation of new thoughts for the context to play a role. In this case, the context is the Background and which parts of it are most effective at the time of a thoughts tokening. This provides a means of understanding not just how some thoughts might be context-sensitive, but how they all might be. This makes it an interesting case to consider for this thesis.

I will then defend this view from three main objections. These are that the view leads to a slippery slope owing to the vastness of the Background, it is hard to know when to stop using its assumptions. The second is that Searle's arguments depend on a contentious analogy between thoughts and natural language sentences. The third is that Searle's view leads to problematic holism.

Chapter 5 discusses Travis's view, in contrast to Fodor's view. I will give an interpretation of Travis's view and contrast it with Fodor's MV. This will provide a contrast between the two and how the arguments they make affect one another. The focus of this chapter will be token-underdeterminacy. I will present Travis's judgement dependent position as a way of resolving token-underdeterminacy without generating further thoughts. The key idea being that when a thought does not have a determinate truth-value we can make a judgement to determine that truth-value. If we judge in

accordance with the way a reasonable judge would, then we succeed in determining the truth-value. What makes Travis a RRC is that he thinks that most thoughts will exhibit token-underdeterminacy as well as most natural language sentences.

This view is very radical, as a judgement by a reasonable judge can determine a thought's truth-value and this is sufficient for the thought to have that truth-value. Whilst this initially seems to get truth backwards, going from judgements to the world rather than using the world to inform one's judgements, this is not the case. This is because being a reasonable judge is difficult and requires that one be sensitive to both the representation and the world. Neglecting the world is not what a reasonable judge does.

In section 1 of this chapter I will introduce Fodor's view, as one that does not allow for token-underdeterminacy at the level of thought. In section 2, I will present some of Travis's arguments against Fodor's view. This section will focus on rule-following arguments to provide some motivation for the Travis-like view that I will introduce. In section 3 I will introduce a judgement dependent view. I will develop this view with a focus on judgement's relation to purpose and action. This helps to bring out the ways in which these judgements are context-sensitive as the purpose for tokening a thought will vary from case to case. In section 4 I will return to two of the key issues in this thesis, the regress argument and the compatibility of judgement dependence with the CTM. I will argue that Travis's judgement dependent view can avoid a regress and be compatible with the CTM.

In chapter 6 I will conclude that weak foundational context-sensitivity is made very plausible by the case of *de se* attitudes that appear to be indexical. So, character-underdeterminacy is a viable form of context-sensitivity at the level of thought. I will also conclude that type-underdeterminacy is possible as a concerning relation seems capable of avoiding a regress. Token-underdeterminacy, understood through the guise of a judgement dependent view, seems to be a viable case of strong foundational context-sensitivity (modal). It avoids the regress argument posed by Fodor and it admits of several context-sensitive factors playing a role in determining the extension of a thought.

In this chapter I will also discuss some of the recurring themes from the thesis. These include the extent to which natural language provides a basis for understanding thoughts and the compatibility of context-sensitive views with the CTM. Whilst both remain controversial it seems that these are future projects once the possibility of context-sensitive thoughts is established.

2. Thoughts

In this chapter I aim to give an account of thoughts for use in the rest of the thesis. The view I set out will mostly follow the language of thought hypothesis (LOTH), initially put forward by Fodor.⁹⁷ This account has some advantages for the approach that I am considering. For one thing, the language of thought (LOT, sometimes called Mentalese) is, by hypothesis, a language. It has a syntax and parts that combine to make whole sentences. This means that we can understand context-sensitivity at the level of thought in a similar way to how we understand it in natural language. For example, it could have indexicals, such as “I”, “you” and so on. That LOT is a language makes it easier to apply models of context-sensitivity from natural language to thoughts. If the conception of thought were one that was not language-like, then it would be more difficult to say whether thoughts were context-sensitive and in what way they might be context-sensitive.

I will begin with an account of the language of thought hypothesis, and the theories in which it is situated. These are the computational theory of mind (CTM) and the representational theory of mind (RTM). I’ll explore one common objection to the CTM – the so-called Frame Problem (see section 2.1.1.2.1) – and argue that, despite claims by some (e.g. Chow 2013 and Samuels 2010), the Frame Problem does constitute a fundamental challenge to CTM. This discussion will lead me to endorse a model of the mind on which there is a non-computational general processor and several computational modules. I will defend modularity from some common objections that have been made about it. I will argue that whilst it seems unlikely that we have a single module that works for all of, for instance, our capacity for vision, it is plausible that there are several modules each concerned with a process in making vision possible.

I will end by providing an overview of the work to be completed in the rest of the thesis considering the accounts of context-sensitivity, from chapter one, and thought, from this chapter. I will then

⁹⁷ See Fodor 1975

highlight a recurrent theme of the work to follow, namely, the role of a regress argument in the claim that context-sensitive thoughts are not possible.

The Language of Thought Hypothesis consists of several claims:

- Thoughts are language like.
- Thoughts represent the world and are useful for explaining many of our actions. This is roughly the RTM.
- Thoughts proceed on a syntactic basis, which is a key claim of the CTM.
- The components of thoughts are concepts.
- These concepts can be composed into sentences that can express propositions.

The point of this chapter is not to argue in defence of this view of thoughts, though I will be arguing that certain assumptions that one needs for this view are defensible. That does not mean that the view is without its strengths. An important point is that the Language of Thought Hypothesis (LOTH) does not have any well-developed rivals that are also able to account for the productivity and systematicity of human thought in a compelling way.⁹⁸ The LOTH is not without its problems. For instance, it is not clear how these concepts are able to refer to the things that they do on a naturalised account. This is not, however, a problem that is unique to the LOTH. The LOTH also has the advantage that classical computation is well understood and does not need there to be something that understands what it is doing or the symbols that it is operating over. So, I will adopt this model even though it has not been conclusively proven, and I will not be proving it here.

2.1 The Language of Thought Hypothesis

The LOTH has been summarised as follows:

⁹⁸ See Fodor and Pylyshyn (2015) Chapter 2, and Cain (2016) chapter 2 for a defence of this view. It appears that even some of the chief competitors of the LOTH, such as connectionism, the view that cognition can be understood as a series of interconnected nodes that stand in complex relationships of activation to one another, seem to be at their most plausible when paired with the Language of Thought.

...mental representations are syntactically structured. Their conditions of semantic evaluation and their causal powers both depend on their syntactic structures; the former because mental representations have a compositional semantics that is sensitive to the syntactic relations among their constituents; the latter because mental processes are *computations* and are thus syntactically driven by definition.⁹⁹

What emerges from this quote is that mental representations are language like. Like natural language sentences, sentences in Mentalese have a syntax, which impacts on their semantics. For example, the sentence “The cat is on the mat” has the semantic content that *the cat is on the mat*, and not the content that *the mat is on the cat*, because of its syntax and the order of its parts. Similarly, the thought THE CAT IS ON THE MAT means that *the cat is on the mat*, and not *the mat is on the cat*, because of its syntax. This means that sentences in Mentalese have a semantics that is determined, at least in part, by the meanings of their parts and their manner of composition.¹⁰⁰ I will assume that, at least, the meanings of sentences in Mentalese must compose to account for systematicity of human thought. I will discuss the compositionality of content below (see section 2.6.2). That these sentences have a semantics means that Mentalese sentences can be understood as a part of a RTM.

The other point from this quote is that Mentalese sentences are operated on in a computational manner. Mentalese sentences are also a part of a CTM. What this means will be discussed below. Whilst it is possible that one has the LOT and RTM without the CTM, this view is unappealing insofar as one then needs to replace the CTM with some other account of how these processes are acted on. This is not an easy task, putting it mildly. So, if one is committed to the LOTH then one is also committed to RTM and, for the moment, to CTM. There are some details that will be important to the account of thoughts that I will give here.

⁹⁹ Fodor 1998: 38

¹⁰⁰ Kuczynski (2007: 228 – 229) seems to argue that CTM cannot get a semantics. I will not address this argument here.

The first is that Mentalese sentences are physically realised in the brain. Fodor, who is most responsible for developing the LOTH, endorses token physicalism, where '[t]oken physicalism is simply the claim that all the events that the sciences talk about are physical events'.¹⁰¹ This leaves it open whether physical tokens also fall under non-physical types. So, there might be a non-physical type, such as expressing the proposition that *the cat is on the mat*, which has a physical token to instantiate it. This non-physical type is individuated according to their functional role, which is a kind of computational role for Fodor.¹⁰² This allows discussion of mental representations whilst also allowing that they are physical events.

Token physicalism is a weak claim. It does not claim that 'every event falls under the laws of some science or other'.¹⁰³ Nor does it claim that 'every *property* mentioned in the laws of any science is a physical property'.¹⁰⁴ This is the allowance for non-physical types to feature in the science, such as representations. Nor does it assume that '...there are natural kind predicates in an ideally completed physics which correspond to each natural kind predicate in any ideally completed special science'.¹⁰⁵ Token physicalism is a non-reductive claim. A given representational type can be realised as one of several different physical types. So, whilst two individuals may have quite different brains, it is still possible that they can think the same thought.

Mentalese sentences are not identical with natural language sentences. Additionally, Mentalese sentences do not need to be represented (though a given thought will still be a representation). They just need to be tokened. So thoughts, as I will understand them, are "naked" in the following sense: '...we must be capable of having naked thoughts ... thoughts devoid of linguistic or other representational clothing: thoughts may represent the world, but, to have thoughts, we need nothing that represents them'.¹⁰⁶ Having a thought just is to have the Mentalese sentence be

¹⁰¹ Fodor 1975: 12

¹⁰² Fodor 2010: 90-91

¹⁰³ Fodor 1975: 12-13

¹⁰⁴ Ibid 13

¹⁰⁵ Ibid 13

¹⁰⁶ Dummett 1996: 166

tokened. To think it, one does not need to have natural language. Having a thought does not mean that that thought we have will be presented to us in natural language as a part of an inner monologue.¹⁰⁷

2.1.1 The Representational Theory of Mind and The Computational Theory of Mind

In this section, I will give an account of RTM and CTM, to make better sense of the LOTH. Both are relevant to understanding the nature of the LOTH, as it often makes use of each to solve various problems that it might encounter. Giving an account of the CTM will require a defence of a modular theory of mind and I will also argue that the scope of computations is limited by what is known as the Frame Problem. I will present the RTM before discussing the Frame Problem.

2.1.1.1 Representational Theory of Mind (RTM)

There is a short and a long version of RTM. The short version can be put as follows ‘Cognitive phenomena are typically the effects of propositional attitudes’ and ‘Relations between minds and propositions are typically mediated by relations between minds and mental representations that express the propositions’.¹⁰⁸ Here we see that mental representations are responsible for our ability to entertain propositions, and these are responsible for some cognitive phenomena.

The longer version breaks down into five theses. I will briefly go over each in turn. ‘First Thesis: *Psychological explanation is nomic and is intentional through and through.* The laws that psychological explanations invoke typically express causal relations among *mental states that are specified under intentional description*; viz. among mental states that are picked out by reference to their contents’.¹⁰⁹ So we can explain X’s actions based on what they believe and desire. These propositional attitudes have aboutness and the aboutness of these claims can be used in

¹⁰⁷ This is a controversial claim, and it excludes some theories that I might otherwise consider. One example is Carruthers’ view, according to which natural ‘...language is the vehicle of non-modular, non-domain-specific, conceptual thinking which integrates the results of modular thinking’.¹⁰⁷ I will not be making use of this position. For discussion see Hermer-Vazquez 2002: 689 and others in the same volume.

¹⁰⁸ Fodor and Pylyshyn 2015: 9-10

¹⁰⁹ Fodor 1998: 7

explanations of an agent's actions. That X's desire was for food explains why they went to the kitchen.

'Second Thesis: "*Mental representations*" are the primitive bearers of intentional content'.¹¹⁰ This means that mental representations do not get their content from natural language, or anything else that's intentional. Rather, it follows that natural language gets its content from Mentalese sentences, or they're both primary. It also implies that Mentalese sentences do not need a further action to determine their content.

That Mentalese sentences do not need further action to determine their content will be significant when considering whether Mentalese sentences can be context-sensitive. It suggests an important disanalogy between natural language sentences and Mentalese sentences. Whist context-sensitivity at the level of natural language could be removed by appealing to a context insensitive Mentalese sentence; this move is not available when considering the context-sensitivity of Mentalese.

As an aside, Fodor's Mentalese view is 'Often targeted on the basis that it holds that mental representations have intrinsic meanings'.¹¹¹ Fodor denies this. Instead, the content of a thought supervenes on its relational properties, on some sort of nomic connection between mental representations and things in the world.¹¹² These connections are taken to be causal relations between the items in the world and the mental representations. The idea here is that '...reference supervenes on a causal chain from a percept to the tokening of a Mentalese symbol by the perceiver'.¹¹³ If so, the mental representation does not have its meaning intrinsically, but simply because it has (the right kind of) a causal relation to the referent. This possibility faces problems when dealing with objects that we have not encountered, or could not encounter, or be in an obvious causal relationship with. It also struggles to distinguish cases in which an object mistakenly

¹¹⁰ Fodor 1998: 7

¹¹¹ Fodor 2005: 109

¹¹² *Ibid* 110

¹¹³ Fodor and Pylyshyn 2016: 86. Things that we cannot see are still grounded in a causal relation, see *ibid* chapter 5.

causes a concept to be tokened, see Chapter 5 for discussion. Attempts have been made to account for these kinds of cases and several others problem cases.¹¹⁴ For the moment I will remain neutral as to whether these solutions are successful. However, it shows a way in which it might be possible to resist the idea that meanings must be had intrinsically.

This suggests a view according to which a Mentalese sentence acquires its content in virtue of causal relations with items in the world. These Mentalese sentences do not require any other kind of representation to get their content. In that sense, they are the bedrock of representation, even though the ability to represent depends on various causal relations.

'Third Thesis: *Thinking is computation*'.¹¹⁵ This ties the RTM to the CTM, and I will elaborate on what it means for thinking to be computational in the following section. For the moment computation is a way to reliably respect semantic properties based on a representation's syntactic properties.

'Fourth Thesis: *Meaning is information (more or less)*'.¹¹⁶ This vague claim is, perhaps, elaborated on in later work by Fodor. He claimed 'that reference is the only semantic property of mental or linguistic representations'.¹¹⁷ So it seems that meaning is just reference on this view. This uses a different notion of propositions to the one that I have advocated in Chapter 1. The notion of proposition that I am making use of is propositions as intensions. So, I will not agree with Fodor on this point. However, a different understanding of the claim that meaning is information allows for the information to be broader than reference alone, it can include intension. For the most part I will continue to take meaning as intensions. However, in Chapter 5 I will consider Fodor's view as it was originally presented by him.

¹¹⁴ Fodor and Pylyshyn 2016: 133ff

¹¹⁵ Fodor 1998: 9

¹¹⁶ Ibid 12

¹¹⁷ Fodor and Pylyshyn 2016: 1

'Fifth Thesis: Whatever distinguishes *coextensive concepts* is ipso facto "*in the head*". This means something like that it's available to be a proximal cause (/effect) of mental processes'.¹¹⁸ This leads Fodor to reject intensions. For example, that Hesperus is Phosphorus (is Venus) is an informative statement because each name as a different intension. The issue is that the reference is the same in each case, which suggests that there is more to the meaning than just the reference.¹¹⁹ Fodor's response is that 'Thoughts and concepts are individuated by their extensions *together with their vehicles*'.¹²⁰ So, "Hesperus is Phosphorus" is different from "Hesperus is Hesperus" because the "vehicle" differs in the first case and not in the second. The token that I use to instantiate it is different, and the information that I gain in learning that two vehicles refer to the same object is what makes the statement informative. However, this response need not exclude other options, such as intensions, taking a role here. It is Fodor's commitment to naturalism that leads him to reject intensions.¹²¹

In summary, the representational theory of mind seeks to explain some psychological phenomena based on the mental representations that people entertain. These representations are the primitive bearers of intentional content and they get content from their relations to their referents. Content can be thought of as being an intension. Concepts are individuated by their extension and the vehicles that are used to instantiate them. They are instantiated in the brain and can be operated over computationally. I will now elaborate on the CTM.

2.1.1.2 The Computational Theory of Mind (CTM)

The CTM has three main tenants, as follows:

- i. Thoughts have their causal roles in virtue of, inter alia, their logical form.

¹¹⁸ Fodor 1998: 15, It should be noted that Fodor is not an internalist about concepts. He maintains that we acquire concepts via a causal link to the extension of that concept (Fodor 2015). He also claims that for an object to be a doorknob, it must be the sort of thing that is disposed to cause us to think DOORKNOB (Fodor 2010: 166).

¹¹⁹ Frege 2010/1892: 217

¹²⁰ Fodor and Pylyshyn 2015: 74

¹²¹ Ibid 3

- ii. The logical form of a thought supervenes on the syntactic form of the corresponding mental representation.
- iii. Mental processes (including, paradigmatically, thinking) are computations, that is, they are operations defined on the syntax of mental representations, and they are reliably truth preserving in indefinitely many cases.¹²²

The notion of a logical form is of central importance here, as it determines what operations are performed on a given representation (alongside the program that is being run). Fodor defines a logical form of a symbol as something that is supposed to make that symbol's '...compositional structure explicit; that is, it's supposed to make explicit the contribution that each of the interpreted parts of the symbol makes to its interpretation'.¹²³ So the logical form of a symbol is something like its grammar made explicit so that it is clear which inferences can be performed on it. So, the logical form of a representation should give us a version of the representation with all its components made explicit. Thus, operations can be performed on a representation which take all its syntactic parts into account.

A key component of CTM is the idea that '...cognitive processes are operations defined over the constituent structures of mental representations of the concepts and propositions that they apply to, which they may supplement, delete, or otherwise rearrange'.¹²⁴ The idea is that thinking is composed of operations on the logical form of a representation. When we have a sentence in Mentalese, this sentence might express a proposition. However, as a computation, the computational device operating on these does not make use of these features. It will only focus on the syntax of the sentence.¹²⁵ The 'mental representations are individuated by... their constituent structure'.¹²⁶

¹²² Fodor 2000: 18-19

¹²³ Fodor 2010: 175

¹²⁴ Fodor and Pylyshyn 2015: 11

¹²⁵ Heil 2004: 108

¹²⁶ Fodor 2010: 62

Computation means that the operations performed are purely formal. ‘A computation... is a formal operation on syntactically structured representations. Accordingly, a mental process, qua computation, is a formal operation on syntactically structured mental representations’.¹²⁷

Computations are performed based on syntactic features of the representation. To use computation to understand thought, one must assume that thoughts themselves have syntactic structure.¹²⁸ This could explain how mental processes can reliably lead from one true thought to another.¹²⁹ For example, logical relations such as *Modus ponens* can be described in a purely formal way.

One advantage of computation is that we can make sense of thinking without the need to postulate a part of the mind which does the understanding for us, which avoids the homunculus problem. This is a problem of trying to understand how we can make sense of intentional content. The challenge is to do so without positing something that understands semantic content (the homunculus). To do so just pushes the question back one step as we now need to understand how the homunculus understands semantic content. In place of a homunculus, we can postulate a “processor” which has set reactions to different syntactic features.¹³⁰ The device appears intelligent in so far as it corresponds to these logical rules in the way that it proceeds from one representation to another.¹³¹ However, it does not understand the semantics of the symbols it operates over. Instead, it has causal, mechanical, reactions to symbols.¹³²

If the processor understood the language, we would face a problem in understanding how the processor understands language, and the danger is that we may have to say that the processor has a processor and so on. It is this sort of regress which the CTM hopes to avoid.¹³³ So, the processor will

¹²⁷ Fodor 2000: 11

¹²⁸ Ibid 13

¹²⁹ Ibid 13

¹³⁰ Pinker 1994: 75

¹³¹ Ibid 76

¹³² Whilst this is an important problem, it is not Fodor’s main motivation for endorsing the CTM. The motivation seems to be largely that the alternatives do not account for various features of representation or our cognitive abilities, such as the systematicity and productivity of thought. For more discussion see Fodor 2000 and 2005 for discussion.

¹³³ Pinker 1994: 75

only be able to operate on the syntactic features of the representations it entertains. It cannot interact with semantic features.

Computation, and its importance, are summed up nicely in the following quote:

The critical property of the machine language of computers is that its formulae can be paired directly with the computationally relevant physical states of the machine in such fashion that the operations the machine performs respect the semantic constraints on formulae in the machine code. Token machine states are, in this sense, interpretable as tokens of the formulae .¹³⁴

According to the CTM, the brain can be in several physical states and some of these states are interpretable as tokens of formulae by a processor that reacts to these formulae in a causal fashion to produce new formulae. The processor respects the semantics of these formulae by respecting the syntax of these formulae. The syntax is determined by the physical state of the brain and so these computational processes can proceed without a need for something that understands what it deals with. Though these states can still be causally relevant.

It should be noted that these computational processes are not considered to be exhaustive of all thought processes. Some thought processes may not be computational on this view. It has been claimed that '...there are parameters of beliefs... that determine their role in nondemonstrative inference but are, on the face of them, not syntactic: relevance, conservatism, simplicity are extremely plausible examples'.¹³⁵ Given that CTM posits computations that are only sensitive to the syntactic features of that representation, it follows that CTM cannot account for processes that rely on these features.

¹³⁴ Fodor 1975: 67

¹³⁵ Fodor 2010: 124

2.1.1.2.1 *Frame Problem*

Properties such as relevance, which do not seem to be traceable to syntax, reveal a problem with a thoroughly computational theory of mind. This is the Frame Problem in artificial intelligence. When it comes to nondemonstrative inferences, e.g. inductive inferences, a decision needs to be made regarding which beliefs to use to make the inference.¹³⁶ We seem to select only those premises that are relevant to our considerations.¹³⁷ For instance, if I want to know why a flower pot has been knocked over in my home, I might consider the premises that I own a cat, that the pot was near an open window and that the cat tends to jump through open windows. I would be unlikely to consider that I was wearing shoes when I left the house. Here some beliefs are selected as relevant to reasoning, whilst others are never considered as a part of that individual's reasoning process. This is owing to their irrelevance.

Given that agents do not make use of all the beliefs available to them, it seems that they must have a means of deciding which of their beliefs to use.¹³⁸ However, this mechanism will not be computational. At least, it cannot be a computation based on the syntactic features of the representation. It cannot be done on a syntactic basis because the relevance of a given belief will vary depending on the hypothesis that one is seeking to prove but the syntax will remain fixed. So, it is hard to see how a computation could be sensitive to relevance. To see this, consider that if I want to know where my shoes are, the belief that I wore my shoes this morning may well be relevant, even though it was irrelevant in the earlier piece of reasoning. This is despite the syntax of these beliefs staying fixed. The problem is that '...not more than a small subset of relevant background beliefs is actually consulted' when making abductive inferences, and this is hard to account for computationally.¹³⁹

¹³⁶ Ibid 116

¹³⁷ Fodor 2000: 37-38, Borg 2004: 82

¹³⁸ Borg 2012: 79

¹³⁹ Fodor 2000: 37

Whilst not all thoughts will proceed one another based on their syntactic form, it seems that all thoughts will still have a syntactic form. This is plausible as we can take conclusions from inductive inferences and use them in making deductive inferences. Inductive and abductive inferences are difficult to account for using syntax because these inferences use features such as relevance of propositions.

This problem leads to a view according to which computations operate on pre-set subsets of information. So, the sub-set of information that the computation must work on is fixed in advance. This has led to the postulation of mental modules, where a module is designed to deal with a particular sort of information. So, it deals with that information and only that information, it does not make use of additional information available to the rest of the mind, 'modules are informationally encapsulated by definition'.¹⁴⁰ So, modules do not have to deal with the Frame Problem. A computation does not select the representations it will operate over. Instead, the selection is a product of the architecture of the mind. This means that not all thought processes will proceed based on syntax and computation alone, so the mind is only partially computational. Non-demonstrative inferences seem to operate non-computationally. How these sorts of inferences are performed is not well understood.¹⁴¹ I will elaborate on mental modules below. First, I will discuss some responses that have been made to the Frame Problem. There are two that are particularly pressing here.

The first line of response is to accept that the Frame Problem is a real problem for computation, but that an advocate of the CTM does not need to worry about it as the human mind does not solve the Frame Problem either.¹⁴² The second accepts that humans do solve the Frame Problem, but argues that classical computation can also solve the Frame Problem.¹⁴³ In which case it is no longer a problem. On either line of response, it is possible to maintain that the mind is thoroughly

¹⁴⁰ Ibid 63

¹⁴¹ Neural networks may be able to account for this if they are a tenable approach in general. See Fodor 2000.

¹⁴² Chow 2013

¹⁴³ Samuels 2010

computational. I will take each in turn and argue that they fail in responding to the Frame Problem. So, we should not have a thoroughly computational theory of mind.

2.1.1.2.2 Chow

Chow argues that the version of the Frame Problem that Fodor endorses is not one that humans solve, at least not in our day to day lives, because it is very demanding.¹⁴⁴ If humans do not solve the Frame Problem, then it hardly matters that computational systems cannot. It does not matter because we want an account of how the mind does what it can and if solving the Frame Problem does not outstrip what computations are capable of then the mind may operate in a computational way. Solving the Frame Problem in this way allows for a thoroughly computational theory of mind, according to which all the mind operates in a computational way. I want to resist this picture of the mind, to allow for a general processor that does not operate computationally. So, I will first give Chow's response to the Frame Problem before responding in favour of the Frame Problem. Whilst the Frame Problem restricts the scope of a CTM it is not fatal to CTM.

Chow distinguishes between several different iterations of the Frame Problem, which he also calls relevance problems. Two will be relevant for the following discussion. The first is what he calls the epistemological relevance problem (ERP), which he states as follows:

Epistemological Relevance Problem: The problem of how a cognitive system considers (mostly) only what is relevant, or how a cognitive system knows what is relevant.¹⁴⁵

This is the version of the Frame Problem that is attributed to Fodor and reflects the idea that it is difficult for a computation to determine relevance for the reasons given above.

The second is the computational relevance problem (CRP), which goes as follows:

¹⁴⁴ Chow 2013: 310

¹⁴⁵ Ibid 314

Computational Relevance Problem: The problem of how a cognitive system tractably delimits (i.e., frames) what gets considered in a given cognitive task.¹⁴⁶

This iteration of the Frame Problem seems similar to Fodor's version, given that the cognitive system in question is to be understood as a computational problem. However, the two are importantly different. The ERP places a stronger requirement on the information that is to be considered. Chow understands this requirement in terms of finding that information which is objectively relevant to the task at hand. Here: 'Objective relevance [...] refers to a kind of relevance that exists independently of cognizers. When x is objectively relevant to y, x bears on y in ways that support certain counterfactual propositions, such as "if x were different, then y would have been different"'.¹⁴⁷ The CRP does not require that cognizers are able to do anything so demanding and this makes it more amenable to being solved by the use of heuristics.¹⁴⁸ Heuristics become useful because there is no special requirement on the information that is sorted. All that is required is that the computation can ignore a proper subset of the information available to it. In which case, the ERP is more demanding. Chow argues that it is so demanding that we, as humans, rarely solve it without exerting a lot of effort. Two reasons are given for this. The first is that determining objective relevance is a difficult task, such that it is normally only a guiding principle in the sciences.¹⁴⁹ The second consideration is that human cognition normally takes a "good enough" approach to reasoning.¹⁵⁰ So, if we can reason well enough for most of our purposes without needing to make use of objectively relevant beliefs, we will. It seems that we can reason well enough for most purposes without using only objectively relevant beliefs. So, humans do not search for objectively relevant beliefs and so do not solve the ERP.

¹⁴⁶ Ibid 314

¹⁴⁷ Ibid 322

¹⁴⁸ Ibid 321

¹⁴⁹ Ibid 323

¹⁵⁰ Ibid 324

This line of argument seems to leave the proponent of the ERP in a dilemma. Either we do not solve the ERP and there is no Frame Problem that cannot be solved by heuristics (e.g. the CRP) or maintain that we do solve the ERP and be committed to an implausible view of cognition. Neither option is palatable if you think the Frame Problem is unsolvable on the CTM.

2.1.1.2.3 Response to Chow

In response, I will argue that this is a false dilemma. Whilst there is a version of the ERP that seems to be subject to these concerns, these problems stem from a commitment to objective relevance and the claim that this is the key notion of relevance that causes a problem. One response, then, is to deny that ERP needs such a strong notion of relevance. A weaker, more subjective, notion of relevance may still be sufficient to get the ERP off the ground. A second response is that Chow misrepresents the Frame Problem by comparing it to the pursuit of knowledge in the sciences.

2.1.1.2.3.1 First Response to Chow

On the first response: it may be that a subjective notion of relevance is sufficient for an ERP problem to occur. A subjective notion of relevance would pick out those things than an agent believes to be objectively relevant. This makes subjective relevance weaker than an objective relevance as the information selected does not have to *be* objectively relevant. It is plausible that an agent will neglect certain pieces of objective relevant information and include pieces of information that are not objectively relevant. In this respect, a subjective notion of relevance seems to be a better fit than objective relevance for the “good enough” approach that human cognition seems to take. This subjective notion of relevance makes the ERP more demanding than the CRP, as the partition of information to be used must still appear to be objectively relevant to the subject.

As the standard for subjective relevance can vary, there will be cases in which something will only count as subjectively relevant when there are many reasons to suppose that it is relevant. This may not always aim to establish objective relevance but to establish that it is reasonable to treat it as relevant under the circumstances. This understanding of relevance is not an arbitrary delimitation of the beliefs to be considered but a process with some consideration towards achieving relevance.

There may also be cases in which the standard is lower still, and less work is needed for a belief to be relevant as in cases where we assume that something is relevant without any serious consideration. (I will also argue that heuristics do not solve the CRP so easily as they are difficult to use on a CTM, see section 2.2.2. If so, the CRP is still not an easy problem to solve).

A worry with this approach is that subjective relevance must still be able to explain how it is that we are able to navigate the world as successfully as we do. If we only use information that we believe to be relevant to a task and we are fallible in determining what is relevant, then we might not reason as successfully as we ought to. One might want to push back on this use of subjective relevance by claiming that what we take to be subjectively relevant is often what *is* relevant. In which case, subjective relativism might still be a demanding notion that is not so distinct from objective relativism. In which case, the ERP is still one that we do not solve even if we have a subjective notion of relevance, because it is still very demanding. It may not be so easy for a position to go between the two horns of Chow's objection without falling onto one or the other.

In response, it seems that Chow could not claim that we operate in so infallible a way (or close to it to be demanding enough to invite a comparison to objective relevance) whilst also claiming that we are making use of a good enough reasoning approach. It may be that what is subjectively relevant is often enough close enough to what is objectively relevant for us to function successfully. So subjective relativism still seems to be importantly distinct from objective relevance.

It is worth adding that the methods by which we determine what is subjectively relevant need not be as demanding as those required for objective relevance. People may easily form beliefs as to what is and is not relevant for a given piece of reasoning without having to engage in a lengthy process of reasoning or engaging in a scientific enquiry. Nevertheless, they do believe these things are relevant. I will elaborate on this point below. Alternatively, one can claim that we are not so fallible, but then objective relativism might be used in a version of the ERP. On either version there is a version of the

ERP that we do seem to solve. Whether we could use heuristics is something that I will discuss in section 2.2.2.

2.1.1.2.3.2 Second Response to Chow

There is a sense in which subjective relevance aspires towards objective relevance. When we form beliefs, it seems that an aim is for our beliefs to be true. So, if I believe that x is relevant for y , I want, or should want, that belief to be true. In following through with this commitment to truth one might be led towards an increasingly demanding process of determining what is true which might lead one to objective relevance. Chow is right that science is both a demanding process and an approach to what is true. However, it seems to be a mistake to think that when determining relevance in day to day life we need something that approaches scientific rigour. This is plausible because humanity was able to exist, presumably making inferences that were not so misguided, before the invention of science. This is the second criticism of Chow, which is that a comparison to science misrepresents the kind of problem that we are solving when we are solving the ERP.

If we understand the ERP in terms of subjective relevance, rather than objective relevance, then it seems far more plausible that this is a problem that we do, in fact, solve. It does not require that we are able to determine what is and is not objectively relevant. This makes the problem a far less demanding one for us to solve, at least when we are not aiming for the kind of rigour as found in the sciences. This allows the ERP to persist even considering the problems that Chow raises, namely, that humans make use of a good enough reasoning approach that does not aim for objective relevance. One solves the Frame Problem when one has chosen those beliefs or sets of information that one believes are relevant for a given task. Whilst one aims for these beliefs to be objectively relevant, it seems that we do not always require them to be to survive. This version of the ERP also seems to pose a problem for the CTM in so far as subjective relevance is not determined by syntax and the subjective ERP seems to be one that we do, in fact, solve.

Chow may respond that we do solve the subjective ERP but claim that this is just a version of the CRP and so does not really undermine his argument. It may be that the subjective ERP is also, for

instance, solved by heuristics which suggests that the two are quite similar. Nevertheless, there is a sense in which the subjective ERP is more demanding than the CRP. The subjective ERP requires that the beliefs sorted are ones that appear relevant to the subject. It is possible to solve the CRP without having solved the subjective ERP. So, it is a mistake to think that the two are the same.

Chow's argument is that either we as human beings do not solve the Frame Problem, so it does not matter that the CTM cannot or maintain that we solve the Frame Problem and be committed to an implausible account of cognition. My response is that we can make use of a weaker notion of relevance than Chow supposes and that in comparing solving the Frame Problem to scientific enquiry Chow misrepresents the Frame Problem. In which case, the Frame Problem remains a challenge to a thoroughly computational theory of mind. In the next section I will discuss Samuels' objection to the Frame Problem.

2.1.1.2.4 Samuels

A second line of response to the Frame Problem comes from Samuels. Whilst Chow argues that there is no gap between human cognition and classical computation because cognition does not solve problems that cognition cannot, Samuels argues the other way around. Samuels accepts that humans do solve the Frame Problem but argues that this does not pose a problem to the CTM because it is possible for computations to solve the Frame Problem as well. Samuels defends a version of the CTM that Fodor has criticised. So much of Samuels' response is a criticism of Fodor's attack on this position. I will offer a response to Samuels' criticisms.

When considering the Frame Problem, which he also calls a relevance problem, Samuels says that 'What is required to turn [the relevance problem] into an objection to CCTR [the classical computational theory of reasoning] is some plausible elaboration of the problem on which it is implausible that CCTR can accommodate the sort of relevance-sensitivity characteristic of human reasoning'.¹⁵¹ Samuels argues that any such elaboration either depends on an implausible account of

¹⁵¹ Samuels 2010: 283

computation (which is what I am referring to as a thoroughly CTM) or human reasoning. Given that Chow makes a similar point about human reasoning, I will focus on Samuels arguments about the implausible nature of computation needed to make the Frame Problem an objection to the CTM. I will outline what Samuels takes to be an implausible account of computation and his reasons for rejecting it. I will then defend that view of computation. Again, this is to maintain that the Frame Problem is a problem for a thorough CTM to motivate a non-computational general processor in the mind. It need not undermine the CTM entirely.

Samuels targets Fodor's view on the nature of computation, in particular he focuses on the E(CTM) which is defined as 'E(CTM): The causal role of a mental representation, R, in cognitive processes is determined by its essential syntactic properties – that is, its *constituent structure*'.¹⁵² This assumption appears critical for Fodor's understanding of the Frame Problem as E(CTM) can be used to claim that relevance is not tracked by computations because relevance is not a syntactic property. Samuels argument against this view is a *modus tollens*, on the basis that the E(CTM) implies that 'R's causal role is wholly determined by R's essential, hence, context invariant properties'.¹⁵³ Samuels argues that this is false by listing factors that determine the causal role of a representation other than its invariant properties. Samuels gives three such factors. These are the program being run, the other beliefs that the agents have and the use of rules and heuristics. Here I will focus on the program being run and other beliefs present, as I will consider heuristics in section 2.2.2.

Samuels points out that 'If the program executed by a computational system is altered, then the causal role of a given representation may well change'.¹⁵⁴ In which case, E(CTM) appears to be false.

¹⁵² Ibid 285. Fodor expresses the E(CTM) slightly differently, beginning with what he calls Principle E, defined as '*Only essential properties of a mental representation can determine its causal role in a mental life*'. I'll use E(CTM) as a name for the doctrine you get when you do read the Computational Theory of Mind as entailing principle E' (Fodor 2000: 24). This leads him to the conclusion that 'Mental processes are sensitive solely to the syntax of mental representations' (Fodor 2000: 24). So, Samuels reading of Fodor seems fair.

¹⁵³ Samuels 2010: 286

¹⁵⁴ Ibid 286

The program is not a part of the representation, yet it changes that representation's causal role. In which case, E(CTM) cannot be used as a premise in Fodor's argument.

The second factor that appears external to a given representation that nevertheless seems to affect the causal role of a representation are the other beliefs available to the process. Samuels claims that 'Even where no change is made to the program, rules or heuristics for a classical system, the mere addition of new representations to a database will typically change what inferences a representation can participate in'.¹⁵⁵ For example, the representation "if I don't want to get rained on, then I should bring an umbrella" won't cause the representation "I should bring an umbrella" to be tokened unless I also have the representation "I don't want to get rained on". This seems to be an example of the causal role of a representation changing depending on factors other than the syntax of that representation. The conditional goes from having no causal role to having the causal role of causing the consequent to be tokened.

In place of the E(CTM), Samuels argues that M(CTM) provides a more plausible account of the CTM. The M(CTM) is the view that '*[t]he role of a mental representation in cognitive processes supervenes on some syntactic facts or other*'.¹⁵⁶ M(CTM) is less restrictive than E(CTM) and allows for syntactic properties other than those inherent to the representation to play a causal role more easily than E(CTM). Samuels also thinks that M(CTM) is not as susceptible to the Frame Problem as the E(CTM). If the M(CTM) is plausible, then it opens alternative ways of solving the Frame Problem and suggests that the Frame Problem is not fatal to a thoroughly computational theory of mind. Fodor is critical of the M(CTM) but Samuels has some arguments to defend it. Here I will present Fodor's criticism of the M(CTM) as a solution to the Frame Problem and Samuel's defence of M(CTM).

2.1.1.2.4.1 Fodor's Criticism Of the M(CTM)

Fodor's criticism of the M(CTM) can be seen as a dilemma. Either the M(CTM) gives up on the idea that the mind operates computationally or the M(CTM) is committed to implausibly large units of

¹⁵⁵ Ibid 286

¹⁵⁶ Fodor 2000: 29

confirmation.¹⁵⁷ I will take each option in turn. On the first horn, Fodor claims that ‘...by definition, which Classical computations apply to a representation is determined not just by some syntactic properties or other but... by its constituent structure, that is, by how the representation is constructed from its parts’.¹⁵⁸ Here Fodor is claiming that if a process uses syntactic properties that do not belong to that particular representation, as suggested by M(CTM), then they are no longer making use of a computational process. Instead, it is just a process that is (sometimes) equivalent to classical computation in terms of its inputs and outputs.¹⁵⁹

The second horn of the dilemma is that if one wants to use the M(CTM) in a way that is classically computational and not only computationally equivalent then one will be committed to having implausibly large units of (dis)confirmation. Suppose that a representation R is in a syntactic relation S to a theory T, but S is not a part of the syntax of R. If M(CTM) is to operate in a classically computational way, then the processor will not be able to act on S. To do that, the relation S would have to be added to R via conjunction. For Fodor, the way to do this is to conjoin R with the relevant parts of T.¹⁶⁰ As any part of T could be relevant to the computational role of R then it seems one will have to conjoin all of T to R to make S computationally readable for any S. The issue for Fodor is then that the entire theory seems to have become a unit of confirmation or disconfirmation when one wants to know whether R is true or not.¹⁶¹

Fodor has two problems with M(CTM). The first is that he finds holism to be intractable given the potential size of the representation.¹⁶² Secondly, Fodor thinks that it should be possible to have units of confirmation smaller than entire theories on what he calls epistemic grounds.¹⁶³ By this he means that ‘[i]t’s also that assessments of confirmation can be, should be, and generally are called for in

¹⁵⁷ Ibid 30-33

¹⁵⁸ Ibid 30

¹⁵⁹ Ibid 30

¹⁶⁰ Ibid 31

¹⁶¹ Ibid 31

¹⁶² Ibid 31

¹⁶³ Ibid 31-32

respect of objects much less elaborate than the totality of one's cognitive commitments'.¹⁶⁴ Here Fodor, is claiming that we should be able to make local inferences and not only holistic ones. So, if the M(CTM) requires that we can only make holistic inferences then it misses this important point and should not be endorsed.

2.1.1.2.4.2 Samuels' Defence of the M(CTM)

I will now discuss Samuels' two main responses to this argument. The first is that Fodor's argument only works if one assumes that we will reason so as to guarantee that we reach the result that we should. Samuels denies that this is something that we do. The second response that Samuels makes is that Fodor confuses units of confirmation with beliefs that are relevant for (dis)confirming a representation.¹⁶⁵ If either of these points goes through then the M(CTM) seems like a plausible account of the CTM and it seems capable of avoiding the Frame Problem. I will discuss each objection in turn.

On the first of these objections, Samuels seems to think that Fodor's understanding of how to account for global properties of representations and inferences is designed to guarantee that these inferences can be done successfully. That is, Fodor's claim that an entire theory be conjoined to a representation has the aim of making sure that the result of this process is guaranteed to be correct (where correctness can be understood as having a theory which is internally consistent and as simple as possible). In which case, the assessment of the theory and R's place in it ought to be a global inference in that it concerns the entire theory. Samuels goes on to claim that:

...it is not enough for Fodor's purposes that such assessments ought to be global. Rather, it needs to be the case that the assessments humans make are, in fact, global; and to my knowledge, there is no reason whatsoever to suppose that this is true. So, whilst it may well be the case that we assess beliefs for their simplicity and conservatism, it's far

¹⁶⁴ Ibid 32

¹⁶⁵ Samuels 2010: 287 – 288

from clear that this means that cognitive processes are sensitive to the global properties of beliefs.¹⁶⁶

Here Samuels is denying that we need to make sense of reasoning according to the M(CTM) by making use of a conjunction of R and T on the basis that this is a model of how we ought to reason but that this is not an accurate model of how we do in fact reason. The argument being that humans are susceptible to more errors in reasoning than this model would suggest. In which case it is possible to make use of a weaker model which does not depend on conjoining R with an entire theory. In which case the M(CTM) need not be committed to holism whilst also being computational.

This response has its problems. One problem is that there needs to be a means by which the M(CTM) can select those representations from T that will be relevant. If it does not simply select all these beliefs, then there will need to be a means by which it selects those that appear to be (most) relevant for the task at hand. At this point we back to square one with Frame Problem.

One option that is available here is something like a consideration generator which, when we are faced with an important decision, will provide us with considerations ‘...some of which may of course be immediately rejected as irrelevant by the agent (consciously or unconsciously)’.¹⁶⁷ Those that aren’t rejected go on to play a role in determining the actions of the agent. Whilst this consideration generator provides a random way of getting premises that a computation could perform, I do not think that it successfully avoids the Frame Problem.¹⁶⁸ Whilst it provides considerations in a way that does not require there to any appreciation of the relevance of those considerations to one’s task, it does not explain how an agent can then decide which considerations are relevant. If this model is going to be used to salvage the thorough CTM there needs to be some account of this step that is computational.

¹⁶⁶ Ibid 287

¹⁶⁷ Dennett 1979: 295

¹⁶⁸ Dennett (1979: 298) points out that the selection of considerations given as output by the consideration generator might be entirely determinate, there need not be any pattern to those that are considered. They need not be ones that have been selected for relevance, for instance.

The second response that Samuels makes is that Fodor's argument confuses units of (dis)confirmation with representations that are relevant to determining (dis)confirmation. He claims that Fodor's argument fails because

...it collapses a distinction between: The units of confirmation: roughly, that which gets confirmed or disconfirmed;

and

Those considerations relevant to assessing or confirming something.

At most what the argument supports is the conclusion that lots of beliefs—that is, K—will need to be accessed by a classical computational device in order that it be sensitive to the global properties of a representation. But that doesn't show that K is part of the unit of confirmation—that is, that which gets (dis)confirmed. All it shows is that K is among those things relevant to the (dis)confirmation of R.¹⁶⁹

Here "K" is used in a similar way to the way that Fodor uses "T". In each case they stand for a set of additional representations or beliefs that are required for determining the causal role of R. Samuels' claim here is that just because K is important for determining whether R is accepted as true, it does not mean that K must be conjoined with R. In which case, one can confirm R is true without having to accept that all of T (or K) is true. One way of seeing the difference between the two cases is to consider two different ways of representing *Modus ponens*.

One way of representing *modus ponens* is as $((P \wedge (P \rightarrow Q)) \rightarrow Q)$. Here it has been understood as a single representation. If one wanted to know whether Q was true on Fodor's view then one would also have to conjoin $(P \wedge (P \rightarrow Q))$, which would stand in for T. This would capture the idea that the units of confirmation are larger representations on M(CTM). On Samuels' understanding of the CTM, taking all these premises and conjoining them into a single representation to prove Q is unnecessary.

¹⁶⁹ Samuels 2010: 287

The premises P and $(P \rightarrow Q)$ can be understood as distinct representations that are relevant to determining Q but need not be conjoined to Q .

Samuels' version of the theory also seems to be a plausible account of how computations might function. So why shouldn't the $M(\text{CTM})$ be understood in this way, instead of the way that Fodor proposes? Either option seems, *prima facie*, to be classically computational and Fodor must want Samuels' version of, for instance, writing *modus ponens*, to be possible in the case of $E(\text{CTM})$ as he does not seem to think that this account leads to holism. It now appears as though $M(\text{CTM})$ can be computational without leading to 'ruinous holism'.¹⁷⁰

2.1.1.2.4.3 Further Criticisms of $M(\text{CTM})$

One response is that $M(\text{CTM})$ is not computational because it does not focus on the local properties of a given representation.¹⁷¹ Which is to say that $M(\text{CTM})$ is not computational because it does not conform to $E(\text{CTM})$. This, however, seems to be what Samuels is questioning. Such a response also runs the risk of making this debate a terminological one, about what it means to call something a computation.

A better response would be to consider what the CTM would look like if it were understood in terms of $M(\text{CTM})$ instead of $E(\text{CTM})$ and consider whether that would be able to avoid the Frame Problem as I have introduced it. Here it appears there is still a problem for the $M(\text{CTM})$. Whilst it may fare better by being able to consider some additional syntactic features of a set representations, it will still have to narrow down the range of representations that it considers for a given task. To do otherwise would, it seems, run the risk of holism or intractability as all representations would have to be considered. There is still a substantial question as to how this narrowing down could be done, how we determine what makes it into K is not obviously solved by using a variety of syntactic features. One possibility is to consider heuristics (see section 2.2.2). Another possibility is to make

¹⁷⁰ Fodor 2000: 33

¹⁷¹ Fodor 2010: 108

use of a modular theory of mind and propose that these modules operate according to M(CTM). It is this possibility that I will now consider.

The key advantage of modules is that they do not require a computation to sort relevant beliefs from irrelevant ones. Any such selection is architectural. So, they would avoid the Frame Problem. On this approach, however, there does not seem to be any benefit to thinking that these modules operate according to M(CTM). In fact, there seems to be a disadvantage. Given that M(CTM) only requires that processes be sensitive to some syntactic facts or other, such a module will need a way of selecting which of these syntactic facts to consider. However, it is unclear how they should be able to do this. Modules do not allow for cognitive penetration, which roughly means that their operations are not affected by cognitive processes that occur outside of that module. In which case, it is unclear how they should receive any instruction in how to perform their task, including the selection of syntactic facts to consider.

By contrast, E(CTM) maintains that only essential properties of a representation, i.e. that representation's syntax, can determine its role.¹⁷² So there is no problem of deciding which syntactic features to use. The M(CTM) computations within a module could, of course, do the same but then there is no point in advocating M(CTM) over E(CTM). The proponent of M(CTM) might say that just as the information available to a module is architectural, so too are the syntactic facts that that module will focus on. That may be so, but it is only a sketch of a solution and needs to explain why, for any given proposal, it is preferable to E(CTM) and how it operates. Here, E(CTM) has an advantage in that it is well understood how a Turing machine may operate.

A conclusion that one may draw from this discussion of the Frame Problem is that if one wants to make use of a CTM then one must either confine those computations to modules or allow for a

¹⁷² Fodor 2000: 24

general processor that uses computations and solves the Frame Problem by use of heuristics. What I have said so far should support this much. I will argue against the use of heuristics in section 2.2.2.

2.1.1.2.5 Summary

To sum up this section so far, computations operate on the syntactic form of a representation. These computations do not consider the semantic properties of the representation. However, the semantic properties are tracked indirectly by tracking the syntactic features of a representation. The idea is that syntactic properties will be fine grained enough to preserve the semantic content. Whenever we have a difference in the semantics of two thoughts, this will be reflected in the logical forms of these thoughts. So, the operations that we perform will differ in each case (even if in each case we perform a *modus ponens* they will be distinguished from each other by the symbols involved in each case).

It is important that the Mentalese 'formulae can be paired directly with the computationally relevant physical states of the machine in such a fashion that the operations the machine performs respect the semantic constraints on formulae in the machine code'.¹⁷³ However, the CTM will not account for all mental processes, but only those which can be exhausted by sensitivity to the syntactic features of the representation alone.

This leads to a very fine-grained account of logical forms where they are finer grained than the surface grammar that we might find in natural language. For instance, where there are ambiguous phrases in natural language, such as the English word "bank", the thoughts we entertain will differ depending on which of the two meanings of bank is the intended one. So, we may have "BANK1" for the financial institution and "BANK2" for the riverbank. This accounts for differences in behaviour, such as going to a financial institution rather than a river side. This shows that computations are sensitive to the concepts that are tokened in a given case. The relation of words to concepts may be one to many and the concepts may divide the world more finely than the words themselves do. I will

¹⁷³ Fodor 1975: 67

give an account of what concepts are like on the CTM after giving a defence of the modular account of mind.


2.2 Modules

Fodor's Frame Problem leads him to a view according to which not all the processes that occur in the mind are computational. A general processing system is needed which does not depend on computational processes to operate. If it did, then it would lead to the Frame Problem. Instead, computation is restricted to mental modules. These mental modules are unlike the general processing system in two important respects: encapsulation and domain specificity. These features allow mental modules to avoid the Frame Problem. In this section I will give an account of modules and will then discuss some of the challenges they face.

To believe that the mind is modular is to believe that '...the mind contains a number of task-specific subsystems that operate in relative independence of one another' and that these subsystems are encapsulated and domain specific.¹⁷⁴ A sub-system is encapsulated when it does not draw on all the information available to the system. Instead, it can only access a proper subset of that information. An example of encapsulation may be found in the case of visual illusions, which persist even when it is known that there is an illusion. An example of such an illusion is the Müller-Lyer illusion.¹⁷⁵ See the below image for an illustration.

¹⁷⁴ Cain 2016: 66

¹⁷⁵ Robbins 2017: S1



In this illusion both lines are of equal length, but without a marker one line appears longer than the other. This illusion persists even when one knows that they are of equal length. That the illusion persists illustrates encapsulation as, if the system were not encapsulated, one would expect the belief that there is an illusion to change the way it looks.

Modules are domain specific in that they only process a specific domain of information (in contrast to the view of modules as bodies of innate information and as a set of constraints on how to cognitively access linguistic material).¹⁷⁶ In other words, ‘...the information available to perform a task depends on which task it is; and the constraints in virtue of which this is so are “architectural”’.¹⁷⁷ This means that it is a part of what it is to be that module that it only deals with that kind of information. This is a part of the reason why modules are often thought to be good candidates for understanding how the mind processes sensory inputs. Common candidates for mental modules include visual and auditory processing. More controversial examples include semantic and syntactic processors. None of these examples are uncontroversial and towards the end of this section I will briefly discuss candidates for mental modules.

¹⁷⁶ Chomsky 1986: 12 – 14, Collins 2017: 227

¹⁷⁷ Fodor 2000: 63

Other properties that have been suggested for modules include:

- Mandatory operation
- Limited Central Accessibility
- Fast processing
- “Shallow outputs”
- Fixed neural architecture
- Characteristic and specific breakdown patterns
- Characteristic ontogenetic pace and sequencing.¹⁷⁸

The operation of modules is sometimes thought to be mandatory in that we cannot choose not to have them functioning whenever they encounter their chosen input. They operate automatically. We have a limited introspective awareness of the processes that occur within modules. This contrasts with our more general reasoning, to which we usually have much more introspective access. The processing of modules is taken to be fast (taking up to half a second to complete a process) and to yield only a shallow output.¹⁷⁹ On one understanding, ‘...the depth of an output seems to be a function of at least two properties: first, how much computation is required to produce it (i.e., shallow means computationally cheap); second, how constrained or specific its informational content is (i.e., shallow means informationally general)’. So, an output is shallow to the extent that it does not require much computational effort and is informationally general, meaning that the output is not specific or detailed.¹⁸⁰ The output is typically simple.

The claim that they have a fixed neural architecture is that we should be able to find some area of the brain in which these modules are hard wired. It is also the claim that these areas of the brain should be dedicated to that task. This is sometimes spoken of as modules being local.

¹⁷⁸ Fodor 1983: 47-101, Robbins 2017: S1

¹⁷⁹ Robbins 2017: S1

¹⁸⁰ Ibid S1

The claim that modules have characteristic and specific breakdown patterns is the claim that these modules can be selectively impaired whilst other faculties are unscathed. That modules have an ontogenetic pace and sequencing is the claim that these modules develop at a fixed rate in the life of individuals with that module. So, people's development in that module should be similar if they are the same age, other things being equal.

Not all of these will necessarily be distinctive of mental modules. Processes in the general system (which is not modular, as it can make use of a wide array of information for a wide range of subjects) might also be fast or local etc. The distinction between modular and non-modular systems is not in how fast they operate, but in the kinds of operations they perform and how they perform them.¹⁸¹ So, whilst modules are computational and general processing is non-computational, this does not thereby tell us whether the modular process operates quickly. It also may not tell us whether the process develops at a standard rate across individuals.¹⁸²

Modularity is also a matter of degree.¹⁸³ I will take the distinctive features of modules to be encapsulation and domain specificity. To be a module is to have the features of encapsulation and domain specificity to an interesting extent. Other features need not necessarily be indicative of a mental module. Encapsulation and domain specificity are required to avoid the Frame Problem. This is because these features restrict the information for the computation to work on in a way that does not make sorting relevant information a computational problem, but rather an architectural one.¹⁸⁴ This means that a computation is not required to select relevant information. Instead, all the information that is available to the module is treated as relevant. The information that is available to the module is determined by its "architecture". It is a part of the nature of that module that it has that kind of information available to it.

¹⁸¹ This is a controversial view, for examples of people who claim that speed is an important part of a modular process see Todd et al (2005), Frankenhuys et al. (2007) and Apperly et al. (2009: 956-7).

¹⁸² Borg 2018: 519

¹⁸³ Fodor 1983: 37

¹⁸⁴ Fodor 2000: 64

This results in an understanding of the CTM on which there are modules which function computationally. These modules are domain specific and encapsulated. If modules also happen to have some of the other properties listed, that is acceptable as well. However, I will not take these other features to be distinctive of modules. There is also a general processing system which does not operate computationally. This enables the CTM to avoid the Frame Problem. The medium in which both systems operate is Mentalese. This allows them to share information.

The original candidates for modules were taken to be sensory systems, such as vision, which Fodor often refers to as 'input systems'.¹⁸⁵ This original understanding of modules is almost certainly too strong. Advocates of modules have had to accept that these systems cannot be a single module, and that if they are modular then they would have to be broken up into smaller processes. This is due to empirical evidence of cognitive penetration of various parts of vision, which I will discuss below. So, vision might be broken up into smaller modules, such as edge detection and depth perception, rather than being a single unified module. The output from these modules might then be processed by either a separate module or the general processing system to produce vision.

The cost of this approach to the CTM is that its success depends on the existence of modules. If there are no modules, then it seems that we are left with only a general processing system. If the Frame Problem is correct, then the general processing system cannot proceed on a computational basis. In this case, we are left with no computational processes. In which case, understanding the LOTM becomes more difficult. Before discussing the case for modules, I will consider some alternatives to this modular understanding of the CTM.

2.2.1 Massive Modularity

A computational alternative to modularity is massive modularity. On this view, the mind is made up entirely of a network of interconnected modules. On the massively modular account '...there is a more or less encapsulated processor for each kind of problem that [the cognitive mind] can solve;

¹⁸⁵ Fodor 1983: 42-3

and, in particular, that there is nothing in the mind that can ask questions about which solution to a problem is “best overall,” that is, best in light of the totality of a creature’s beliefs and utilities’.¹⁸⁶ As described here, there is no general processing system in the mind for processes like reasoning or judgement. However, a view is still called massively modular even if it only maintains that most of the mind is modular.

The problems with a massively modular account of mind are (i) it becomes less clear what a module is and (ii) it is subject to the Frame Problem.¹⁸⁷ To see these, consider Carruthers’ view of massive modularity. On his view, several features that initially appeared to be important aspects of modularity are dismissed as features of modularity. In particular ‘...encapsulation will likely have to be struck out’.¹⁸⁸ Instead, Carruthers opts for a distinction between narrow and wide scope encapsulation. Narrow encapsulation is what Fodor takes encapsulation to be. A narrow-scope encapsulated system can’t be affected by most of the information that that mind has available in its processing. A wide-scope encapsulated system can’t be affected by *most* of the information held in the mind during its processing but can still make use of some of the information in the mind.¹⁸⁹ In the wide-scope case, the module can make use of frugal search heuristics to select some information for use in the mind.¹⁹⁰

However, the wide scope understanding of encapsulation still seems to require that some information be selected as relevant. This leads to the Frame Problem again. In response, Carruthers, and others, have suggested that the use of heuristics to select information might avoid the Frame Problem. Prinz suggests that search engines like Google form an existence proof of computational

¹⁸⁶ Fodor 2000: 64. Fodor was not a proponent of this view, but his summary of the view is still useful to get a grasp on what the minimal claim of massive modularity is. A key proponent of this view is Carruthers (2006). Others define massive modularity similarly as, for instance, Barret and Kurzban (2006: 628) define massive modularity as the view that ‘...many or most information-processing systems in the mind might be modular as well’. Cosmides and Tooby (1994) also argue for massive modularity based on evolutionary grounds.

¹⁸⁷ Fodor 2000, Robbins 2017

¹⁸⁸ Carruthers 2006: 12

¹⁸⁹ Ibid 58

¹⁹⁰ Ibid 59

systems that can use heuristics to select relevant information for a task very quickly whilst also using heuristics.¹⁹¹ In which case, searching large amounts of information in a short time can be done computationally. Not only that, but the information is then put into order of relevance. In which case, the Frame Problem might be less of a concern than Fodor thought. This is a strategy that both Chow and Samuels were also keen to employ.

2.2.2 Heuristics

However, the use of heuristics may be circular. One concern is that deciding which heuristics to use would be difficult on a computational account for the same reasons that making an abductive inference is difficult on a CTM.¹⁹² A computational process would still have to be able to check for relevant beliefs and decide what role those beliefs would play a role in solving the problem.

Furthermore, deciding which heuristic to use would also be a non-demonstrative inference and so subject to the same problems as abductive inference. For instance, one needs to describe the situation that one is in to find and use the relevant heuristic. This may be important if one thinks that the mind may work like a search engine, as what one searches for will determine the results that one gets. This, however, 'depends on what I'm to take to be the *relevant* description' of the situation.¹⁹³ In which case, the Frame Problem recurs. So, appealing to heuristics does not seem to escape the Frame Problem. Instead, it seems to encounter it all over again.

Fodor's response assumes that finding the correct description of the problem you are facing can only be done in a top-down way, from general processing through to the heuristic. A heuristic process that was not top-down might be able to select an appropriate heuristic without the need for the top-down process Fodor assumes is necessary. Such a process might be based on the perceptual information available, the domain in which the agent is situated, the amount of information available to them, or what they associate with that information.¹⁹⁴ Cue ordering also seems to be a possibility

¹⁹¹ Prinz 2006: 33

¹⁹² Fodor 2000: 42

¹⁹³ Fodor 2010: 119

¹⁹⁴ See Goldstein et al 2002 (184 - 6).

in this regard.¹⁹⁵ Cues are a factor that can be used to determine the output of a decision such as which heuristic to use, and may include some of the factors used above. Cue ordering is the proposal that these cues can be put into an order that produces accurate results in a frugal way by selecting a heuristic that interacts with that cue. A draw-back of this kind of approach is, however, that it is not clear how to select heuristics for novel situations. In this case, however, a random selection of a heuristic would be a bottom up process that could work. In the event of that heuristic failing then an alternative heuristic could be selected until the agent either gives up or a desired output is given.

However, some research suggests that this is not how we select heuristics. Instead, the selection of heuristics is subject to a wide range of factors. For instance, it has been found that factors such as a person's intelligence, the time that they have to make their decisions, how neurotic they are and where they are from can all make a difference to the heuristic that they employ in a given situation.¹⁹⁶ In which case, we use our understanding of the circumstances and the kind of solution we are interested in to select a heuristic. That in turn requires an understanding of relevance to have an appropriate description of that task.

Furthermore, these procedures, such as cue ordering, require a feedback loop that adjusts the ordering of cues based on whether the use of a heuristic led to success or not. This requires some procedure for determining the success or failure of a heuristic's output and this will also vary with the context. In which case, some understanding of relevance seems to be required and it is still not obvious how this could be done computationally.

Whilst debate on the topic of massive modularity is ongoing, I will not be making use of such a theory. This is partly because of some of the reasons that the theory has remained controversial, as mentioned briefly above and partly because massive modularity sits less comfortably with the possibility of Really Radical Contextualism (RRC), according to which nearly all natural language

¹⁹⁵ See Todd and Dieckmann (2004/2014) for discussion.

¹⁹⁶ Del Campo et al (2016).

sentences and thoughts are context-sensitive. A view that does not allow for a general processing system would be difficult to reconcile with the RRC views that I will examine in later chapters. As many factors might be relevant to determining content, having a process that is not encapsulated would be beneficial, especially if this process is still compatible with the CTM. So, if a RRC is going to make use of Mentalese then it seems that a view which admits of some element of general processing is required as the process may not be computational.¹⁹⁷

2.3 Defence of a Modular Account

The existence of modules has been threatened on at least three fronts. These are the threat from cognitive penetration, the claim that domain specificity is not well enough defined to be useful for research (Prinz 2006) and third, that modules are not useful in cognitive science. I will consider these objections in turn. I will suggest that the responses are sufficient to allow this framework to persist and therefore to be used in framing questions about the context-sensitivity of thoughts. This is important as the Frame Problem already suggests that a thorough CTM cannot be correct. Modules allow for a weaker version of the CTM to be preserved.

2.3.1 Cognitive Penetration of Modules

First, there is the threat from cognitive penetrability, roughly defined as follows: ‘A perceptual system is cognitively penetrable if and only if its operations are directly causally sensitive to the agent’s beliefs, desires, intentions, or other nonperceptual states’.¹⁹⁸ The mental state of the agent must play a direct role to qualify as cognitive penetration. This is to rule out indirect cases such as desiring food leading one to go to the fridge and to look inside of it.¹⁹⁹ Here desiring food causes one to see food in the fridge, but that is not a case of cognitive penetration.

¹⁹⁷ The main alternative to RTM is connectionism. Connectionism is the idea that we should understand cognition in terms of patterns of activation between nodes. Unfortunately, I won’t be considering this as it would take me too far from what this thesis concerns; for arguments against connectionist approaches, see Fodor and Pylyshyn (1988). See Symons and Calvo (2014) for an overview of this debate.

¹⁹⁸ Robbins 2007: S2.1

¹⁹⁹ Stokes 2013: 655

Cognitive penetration suggests that a system is not informationally encapsulated. If a system is not informationally encapsulated, then it seems that it is not a mental module. If it were a module, then it is less clear what benefits modules would bring as they could be subject to Frame Problem in selecting relevant beliefs. So, if a candidate for modularity can be shown to be cognitively penetrable then it is either not a module or there is little motivation to advocate for the existence of modules. If all candidates for modularity are penetrable then there is good reason to reject a modular account of mind. Here I will discuss some of the evidence for thinking that some candidates for modules are penetrable and consider some responses. I will argue that there do seem to be some cognitively impenetrable systems and therefore modules. However, these modules are only parts of a broader system. For example, vision seems to be cognitively penetrable but edge detection in vision does not.

It is worth clarifying that cognitive penetrability does not entail that we have conscious control over penetrability. Alleged examples of cognitive penetrability, such as desirable objects appearing closer than they are, are not under conscious control.²⁰⁰ It is sufficient for cognitive penetrability that there be information that is external to the module that can impact on the workings of that module.

Prinz brings some criticisms to bear on encapsulation by offering examples of cognitive penetration and undermining examples that would otherwise support encapsulation by offering competing explanations. I will discuss each in turn.

2.3.1.1 Prinz on Cognitive Penetration

One of the arguments for encapsulation is the persistence of illusions even when we are aware that they are illusions. Proponents of modules use this persistence to argue that vision is encapsulated. If it was not encapsulated, then we should expect the knowledge that we are seeing an illusion to keep us from seeing it. However, it may be that whenever we have a conflict between what we see and what we believe, what we see will trump belief.²⁰¹ So, it needn't be that vision is encapsulated.

²⁰⁰ Stokes 2013: 656, findings originally in Balcetis and Dunning 2006.

²⁰¹ Prinz 2006: 31

It might be objected that there is not really any alternative suggested here. For vision to trump belief in any significant way, it may be the case that vision is not penetrated by beliefs anyway. This would be the most obvious way of explaining why visual input trumps belief. If vision were penetrated by our beliefs, then vision may not be able to trump belief. The output of belief would already be affected by the beliefs. That is, visual processing must be free from the influence of our beliefs at some level. If the two were integrated from the start, then it is not clear how the distinction could be drawn at all. This suggests that there may be some distinction between vision and belief, at least on some level. So, the trumping story may presuppose that vision is encapsulated, at least to a certain extent.

More seriously there is the worry that there are examples of cognitive penetration of sensory processes by our beliefs. There seem to be many such examples. When one is expecting a visitor many noises start to sound like knocks at the door.²⁰² There is also evidence that visual experience can distort auditory experience, as in the McGurk effect, where an auditory stimulus of one sound is played with a visual stimulus of someone pronouncing a different sound leads the subject to hear a third sound.²⁰³ A classic example of this effect is the syllable [ba] being dubbed over utterances of [ga], so that there is a conflict between the auditory information, [ba] conflicting with the visual information of lips making a [ga] sound. However, many normal adults report hearing [da]. Sound can cause people to feel sensory illusions, as when hearing multiple tones leads people to feel multiple taps when there was only one tap.²⁰⁴ In this experiment, subjects had a pin that would lightly touch them on the finger, providing a tactile sensation. Subjects were played a tone by some speakers. Subjects reported having more tactile stimuli than they were subjected to when two or more tones were played.²⁰⁵ Synaesthesia also seems to show that people's senses can cross

²⁰² Ibid 31

²⁰³ McGurk and MacDonald 1976: 746

²⁰⁴ Hötting and Röder 2004: 61

²⁰⁵ Ibid 61-62. Strangely, the effect was reduced when four tones were played.

boundaries.²⁰⁶ In each of these cases, a sensory process seems to be susceptible to information from other senses, which suggests that each process can be penetrated by others.

2.3.1.2 Responses to Cognitive Penetration

A standard reply to these concerns is to provide an alternative explanation for the results that does not depend on cognitive penetration. One response is to point out that:

...the cognitive penetration of *some* components of perceptual processing does not, by itself, imply the cognitive penetration of experience, since conscious experience may be the result of or be identified with some broader class of processing, and certain subsets of perceptual processing may not result in conscious experience at all. For the same reason, one cannot infer from the apparent penetration of experience to the penetration of any particular stage in perceptual computation. However, the cognitive penetration of experience implies the cognitive penetration of perceptual processing *at some stage*.²⁰⁷

The proponent of the modular theory of mind would have to insist that these apparent cases of cognitive penetration occur outside of the modules that are responsible for vision, with the effect instead being the result of processing in the general processing system. (They can also claim that these can be cases of penetration on the lower end of the module spectrum. If this was all there was to say for modularity, then it becomes a less useful concept).

This response seems implausible for two reasons. One is that we do not seem to have any control over whether we are subject to these illusions. Yet we seem to have more control over our general processing system than this suggests. We seem to have some control over what we reason about, at

²⁰⁶ Prinz 2006: 32. Though these subjects are not typical of the general population, with at most 4% displaying some sort of synaesthesia (Marks 2017: 24), and so might not be the best indication of how cognition works for the general population. It is also not clear how best to understand synaesthesia given its complexity and not all possible versions seem to support Prinz's point here (See Marks 2017 for possible understandings of Synaesthesia).

²⁰⁷ Stokes 2013: 653

least, we seem to have more control over what we reason about than we do over what we see or hear.²⁰⁸


The second problem with this defence of modularity is that as examples of what appears to be cognitive penetration pile up there is a risk of making the modular theory of mind look trivial or unfalsifiable. If there are no counter examples, then the theory seems unfalsifiable or forced to have to shrink to smaller and smaller modules as cognitive penetration seems to run deeper than was initially supposed. What is needed is evidence that there are encapsulated modules, even if these are not as wide ranging as was supposed. For instance, that it is not the whole of vision that is a module, but maybe that parts of vision are modular.

2.3.1.2.1 Contour Interpolation

However, the empirical case for modularity is not so bleak as these examples suggest and there is some evidence for the existence of these smaller scale modules. In which case, cognitive penetration only goes so deep and the risk of modularity being trivial appears less pressing. Keane provides a case study of contour interpolation, the process by which the perceptual system represents non-visible edges based on how the surrounding visible edges are spatiotemporally configured.²⁰⁹ An example is the four “Pac-Man” shapes which appear to have a square covering them (see figure below).

²⁰⁸ It is possible to avoid the claim that we have control over what we think by adopting a view of thoughts on which thoughts are merely things that happen to us. A view like this is given by Strawson (2003:236) who argues that thoughts cannot be actions as that would lead to a paradox. We would have already had to have had the thought to consider whether to think it. However, even on this view we have a level of control that distinguishes thoughts from cases of perception. Even on Strawson’s view, we can focus our attention on topics that are not in our immediate environment (Strawson 2003: 231). We cannot do the same with what we see. So, even on a view on which we have little control over what we think, there still seems to be an important difference between our control over what we see and what we think. Note that if Strawson’s view is incorrect and one takes a stronger view over the extent to which thoughts are actions, such as Peacocke’s according to which ‘... the nature of belief, judgement and intentional content cannot be properly elucidated unless we recognise that judgements are actions’, the contrast between seeing and thinking gets greater (Peacocke 1999: 20).

²⁰⁹ Keane 2018: 1



This is a useful example of what might qualify as a smaller subsystem. This system is domain specific as it ‘...determines whether spatially segregated edges belong together and also how they combine to form a continuous edge’.²¹⁰

Keane argues that this is a module as it exhibits all the features that Fodor attributed to modules. Here, what is of importance is encapsulation. Keane points to studies that seem to show that subjects’ perception of objects that have been obscured does not change even when they have been given a prior belief as to what that object will be.²¹¹ This includes cases of priming and cases in which we have prior beliefs about, for instance, the length of a person’s arms. In neither case do these beliefs or priming effects seem to have an effect. This suggests that even when subjects know they are seeing an illusion they are still susceptible to it, suggesting that there is not cognitive penetration of this module.

Apparent cases of cognitive penetration here, that do seem to make use of concepts and belief, are instances of contour abstraction.²¹² In these cases, one makes use of a pre-existing concept that appears similar to the stimulus to fill in the blanks. Interpolation, by contrast, works the other way

²¹⁰ Ibid 3

²¹¹ Ibid 6

²¹² Ibid 7

around. Missing boundaries are filled in which might then be recognised as something familiar. An example of contour abstraction occurs when one sees the head of a dog and an opaque screen in front of where we'd expect the dog's body to be, people can form a reasonable idea of the outline of the rest of the dog. Such cases are susceptible to cognitive penetration and so do not seem to be encapsulated.

There is evidence that contour abstraction and interpolation are distinct processes.²¹³ Contour abstraction is the representation of non-visible edges based on what one believes the visible edges to be edges of. Contour interpolation is the representation of non-visible edges based on how surrounding edges are configured.²¹⁴ Contour abstraction seems to be more susceptible to the will of the perceiver as ambiguous shapes can be filled in in different ways depending on the concept chosen by the subject. So, these cases of contour abstraction are not part of a module. However, this distinction means that the cases such as the dog's head visible over a screen are not cases of contour interpolation. As it is contour interpolation that is claimed to be modular, those examples are not counter examples to the claim that contour interpolation is a modular process. This suggests that there are encapsulated mental processes that are suitable candidates for modularity.

2.3.2 Domain Specificity

A second concern is with the notion of "domain specificity". Prinz has argued that this term is hopelessly vague.²¹⁵ He considers each term individually. On the term "domain" he claims that there is a weak and a strong reading. On a weak reading, it is just to have a subject matter.²¹⁶ But this is too weak and encompasses too much. On the strong reading "domain" does not refer to 'any subject matter, but to matters that are relatively encompassing'.²¹⁷ The objection that Prinz makes is that it is not at all clear what it should be relative to. The concept CAMEL is less encompassing than ANIMAL

²¹³ Ibid 7-8

²¹⁴ Ibid 2

²¹⁵ Prinz 2006: 27 – 28

²¹⁶ Ibid 28

²¹⁷ Ibid 28, though advocates of modules may deny that this is a fair interpretation of their views, it does put some pressure on them to define domain specificity.

but is more encompassing than THE PARTICULAR ANIMAL USED BY LAWRENCE TO CROSS THE DESERT.²¹⁸ This makes the notion of a domain less useful when determining whether we are using a module and makes the notion of modularity itself too poorly defined to be useful. It is possible that any process could be described as domain specific, if Prinz's objections are correct. In which case there is no principled way to say what is and is not a module.

On the term "specificity" Prinz argues that there is also both a stronger and weaker understanding of this term. On the weak reading specificity would mean that '... [a mental resource] is used to process information underlying our aptitude for that domain'.²¹⁹ This reading is too weak as it entails that we have a domain specific capacity to do something in virtue of having the mental resources to do so. The trouble is that these resources could be used to perform several other notably different tasks.²²⁰ This would fail to capture what is intended by "domain specificity".

On the stronger reading, "specific" can mean "'exclusively dedicated'".²²¹ Whilst this concept is more rigorous, Prinz argues that it is not a feature that examples of modules have. For instance, Prinz argues that the neurons responsible for our language capacities are also useful in other domains. For instance, the ability to recognise patterns is important for language but is also used for other potential domains outside of our language capacities.²²² In which case, they are not exclusively dedicated to that task. So, Prinz's objection is that either the notion of specificity is too weak to be useful in carving up the mind or it is too strong and does not apply to any of our actual cognitive processes.

One response to these objections would be to accept that modules are more limited than we had previously thought. For instance, edge detection still seems to be modular, despite its being only a

²¹⁸ Ibid 28

²¹⁹ Ibid 28

²²⁰ Ibid 29

²²¹ Ibid 28

²²² Ibid 29

component of vision.²²³ That alone might prove to be sufficient to maintain the weaker understanding of the CTM.

A response is that we are still able to identify modules, even if they are less encompassing than was originally supposed. It is noteworthy that many of the examples of modularity that Prinz considers are quite coarse, encompassing, for instance, our entire capacity for vision or for language and so on. Taking a finer grained approach to modularity, on which modules deal with smaller domains, still seems to be viable. Keane's work also supports this. So, whilst domain specificity is hard to define it does not mean that it is a useless notion. Rather, whether there is a specific domain or not often seems to be determined in part by the empirical research on cognition.

2.3.3 No Explanatory Work

The third concern is that modules do not do any explanatory work. If modules aren't distinguished from general processing in virtue of their speed, genetic development, shallow output, being automatic and so on, then we might wonder whether there is any point in positing modules. If the general processing component of the mind can do what modules were traditionally thought to do, then perhaps we do not have enough motivation to think that there are modules. The idea of modules seems to be at risk of being cut out by Occam's razor.

If so, we can allow a general processing system to do the work that modules had previously been taken to perform. In which case, there is no reason to think that there are modules. If this is correct, and the Frame Problem is also a serious problem, then it seems that we may lose the CTM. I will elaborate on the worry before responding to the objection.

Even if the CTM were to be lost like this, it might be that RTM would still be preservable minus the thesis that thinking is computation. However, this is not much of a consolation prize, as there is no obvious way to implement RTM without computation. It is less clear what the representations are and how they can have the effects that they do in determining our actions and our reasoning. One of

²²³ Ibid 29

the appealing things about the CTM was that it was able to explain the causal properties of representations without making use of anything that understood the representations. It would still be possible to consider whether thoughts, understood in these terms, are context-sensitive. However, it is much less obvious that these are the terms that we should be using. For instance, if thoughts are not language like then applying context-sensitivity to them in a way that is intended to parallel natural language context-sensitivity will not make sense.

In response to this Occam's razor worry, one might think that the evidence of encapsulation that I have considered above goes some way to showing that there is still a point in postulating mental modules. It can help to make sense of illusions that we are subject to even when we know better and it seems that modules are defensible from other objections that have been made against them. It is also significant that we have not got much of an understanding of how general processing might work. So, opting for general processing as the only thing that the mind does seems to be akin to giving up on understanding how the mind works. Although the jury is still out on modules, it is not obvious that modules have clearly been refuted or confirmed one way or another. However, for as long as it is possible that there are encapsulated and domain specific cognitive processes then it continues to be possible that there are modules.

2.3.3.1 Sternberg on Modular Explanations

Sternberg offers some experimental evidence which suggests that there are several different modules at play in the human mind. Some tasks show differences in either an individual's reaction time or that their brain scans differ in interesting ways. Here I will consider one of the examples that Sternberg considers. In this experiment, subjects were tasked with classifying visually displayed numbers as either greater or less than 65. There were two variables that were manipulated. These were the notation (N) used to present the numbers, which was either in Arabic numerals or in number names. The other variable was numerical proximity (P) to 65.²²⁴

²²⁴ Sternberg 2011: 175

Several experiments have been conducted using this format, all measuring the reaction time (RT) of subjects. Considering reaction time it was shown that there are '... additive effects of N[otation] and P[roximity] on RT; this was interpreted to indicate two modular subprocesses arranged as stages: encoding (E), influenced by N, which determines the identity of the stimulus and is slower for number names than numbers, and comparison (C), influenced by P'.²²⁵ This result suggests that there are processes that are modular which have specific inputs for their domain, and this provides evidence for domain specificity which is an important part of modularity.

This experiment alone does not show that these processes are informationally encapsulated and whilst they do provide evidence that the processes are domain specific this is not proof that they are. Though Sternberg does discuss other experiments that also suggest that there are processes that appear to be modular.²²⁶ Furthermore, if one should take the Frame Problem seriously and consider a general account of intelligence out of reach of a CTM, then rejecting modules leaves you without a way of accounting for any process computationally. As general processing seems to be all that is left to this view, it is left with no account of how any of the mind works at all.

2.3.3.2 Other Cases of Modular Explanation

Furthermore, some cases are difficult to account for without making use of a modular account of the mind. For instance, there are cases in which a person's language capability is deficient in some respect whilst their general intelligence remains normal or vice versa. This runs counter to what one would expect if general intelligence were responsible for everything. If that were the case, then one would expect one's general intelligence to determine one's abilities in other areas. However, that is not the case. Here are some examples.

One is the case of Genie who was deprived of linguistic input from 20 months old to past 13 years old. Whilst her linguistic abilities did not recover as she was, for instance, never able to use

²²⁵ Ibid 176-177

²²⁶ See Sternberg 2011

pronouns, she developed normally in other respects.²²⁷ An example of the opposite effect, good language ability and poor general intelligence, is found in Christopher. Christopher has a low non-verbal IQ, yet he can translate texts from 15 other languages into English at normal reading speed.²²⁸ Again, if humans' language capability were not modular then one would expect Christopher's language capability to poor and not as exceptional as it is.²²⁹

2.3.3.3 Summary

To conclude this section, I have presented a modular account of the CTM. On this view, there are mental modules which deal with a domain of information and are informationally encapsulated. These modules operate according to Turing style computations and use a language of thought. In contrast to these modules, there is also a general processing system that can perform non-demonstrative inferences, and which is not computational. Both systems make use of Mentalese to operate.

I have offered a defence of this view from some objections, including the idea that there are no encapsulated modules, that there are no domain specific modules and that these modules are not useful in our understanding of the mind. Whilst there is still a lot of empirical work to do in settling the dispute on modules, I have argued that the theory is not without use. For instance, though vision as a whole may not be modular, it seems to have parts that are modular. In the next section I turn to giving an account of concepts as the components of Mentalese sentences.

2.4 Concepts

Concepts are what the sentences in Mentalese are composed out of. It is claimed that '...concepts are constituents of mental states. Thus, for example, believing that *cats are animals* is a paradigmatic

²²⁷ MacSwan and Rolstad 2005: 232-233

²²⁸ Ibid 235. Other cases are discussed in this article by I am only going to include these two for the sake of brevity.

²²⁹ See Harris 2018 for other arguments in favour of the impenetrability of linguistic competencies. Whilst there may be alternative explanations of this data, for instance a non-modular general processor with impairments to specific skills might also account for this. It is, however, hard to understand a general processor can be limited regarding such specific faculties unless that faculty was already divided in some way. At which point we seem to be back on the road towards modularity.

mental state, and the concept ANIMAL is a constituent of the belief that *cats are animals*'.²³⁰ So concepts are comparable to words in natural language sentences, in so far as words are the components of sentences. This comparison will only go so far as it may be the case that certain natural language sentences do not have truth conditions independently of their relation to a mentalese sentence. Each concept will have its own contribution to the logical forms that it can be a part of.

In virtue of being a part of both RTM and CTM, concepts will have five features. These are:

1. Concepts are mental particulars.
2. Concepts are categories.
3. Mental representations are compositional.
4. Some concepts are learned.
5. Concepts are public.

I will elaborate briefly on each of these in turn. The point is to give an account of Fodor's view of concepts. However, not all these features will be assumed to be features of concepts going forwards in examining the thesis. For instance, that mental representations are compositional will not be assumed going forwards. This is because Really Radical Contextualism (RRC) is incompatible with this kind of composition, according to which the content of the whole is determined only by the content of the parts and the way they are combined. (It is possible that composition happens after modulation, which is compatible with Fodor's view. I will discuss this in Chapter 5).

'1. Concepts are mental particulars; specifically, they satisfy whatever ontological conditions have to be met by things that function as mental causes and effects'.²³¹ This seems to commit one to the view that mental state types can be are physically realised in the brain as tokens of those concept types. That is, if one thinks that only physical events can have a causal impact on the physical world,

²³⁰ Fodor 1998: 6

²³¹ Ibid 24

and that thoughts have a causal effect on the physical world, then thoughts must be physically realised. This means that thoughts, and the concept tokens out of which they are composed, are physically realised. This brings us back to the commitment to token physicalism. So, tokens of a concept type are physically realised mental particulars.

'2. Concepts are categories and are routinely employed as such'.²³² This is the claim that concepts have extensions. When considering individual concepts, rather than sentences in the LOT, the extension is just the objects which fall under that concept. For instance, the concept CAT has as its extension all the things that are cats. In which case "x is a cat" gives the value true only if x is a cat.

'3. Compositionality: concepts are the constituents of thoughts and in indefinitely many cases, of one another. Mental representations inherit their contents from the contents of their constituents'.²³³ This means that thoughts are composed out of concepts and follow a principle of composition. It seems that, on Fodor's view, the concepts in a Mentalese sentence and their manner of composition is all that there is to be determining the content of that sentence. He has claimed that '[t]he compositionality of mental representations... thus mirrors the compositionality of the propositions they express', where he seems to take propositions to be composed.²³⁴ Fodor has also denied that there is metaphysical underdeterminacy to mental representations, which suggests that concepts and their composition is all there is to determining content.²³⁵

Composition will be challenged in this thesis, so I will not assume that this is correct. However, it seems compatible with context-sensitive views that the concepts are still important components of the thought. It only means that composition is not all that there is to determining the content of a thought. Going forward I will assume that this weaker understanding of concepts constituting

²³² Ibid 24

²³³ Ibid 25

²³⁴ Fodor and Pylyshyn 2016: 9. Here Fodor is using a referential or Russellian notion of a proposition. Fodor makes similar claims about the composition of mental representations elsewhere, see Fodor (2001: 14) and Fodor (1998): 27.

²³⁵ Fodor 2005: 104

thoughts is possible. Nor does a contextualist need to deny that any composition is possible. For example, the meaning, or character, might compose even when the content doesn't (see chapter 1 for more detail on these terms). This is a form of composition that radical contextualists are still willing to endorse.²³⁶

'4. Quite a lot of concepts turn out to be learned'.²³⁷ The concepts that are learned are taken to be complex concepts, and they are complex in so far as they are composed out of other concepts. For instance, a complex concept might be BROWN COW as it seems to be composed out of the concepts BROWN and COW. Basic concepts are taken to be unlearned concepts. This follows from the Representational Theory of Mind (RTM). If one is going to represent things to oneself, one needs to have some concepts in place with which to compose the representation. This will include representations which we will use to learn the content of other concepts. When we learn a concept, we will need a belief about that concept's identity which will require concepts to form that belief. 'learning a concept involves acquiring a belief about its identity'.²³⁸ This presents a problem for the possibility of concept acquisition and is one of the key motivations for Fodor's nativism. The problem is that to acquire a concept by learning it one would need to form a hypothesis about the identity of the concept.²³⁹ So, for a concept G to be learnt one would need a hypothesis about its identity, such as G is F. This requires a prior understanding of that concept, in this case F. In which case, the concept was already known. On Fodor's view, there is no alternative to nativism.

The idea means that we must start with some concepts for this process of concept learning to get off the ground.²⁴⁰ We do not learn these basic concepts; they are in some sense innate. For Fodor, we have an innate disposition to acquire certain concepts because of encountering certain stimuli in the world.²⁴¹ So, the mechanism by which we acquire concepts is innate. As such we acquire stereotypes

²³⁶ Travis 1997, Carston 2002

²³⁷ Fodor 1998: 27

²³⁸ Fodor 2010: 141

²³⁹ Fodor 1975: 35

²⁴⁰ Fodor 2010: 141

²⁴¹ *Ibid* 166

of what a typical instance of a type might be from our experience of the world and these allow us to activate the concepts that were in our repertoire.²⁴² Stereotype formation might be a stage in the process of acquiring concepts.

We have this store of concepts in virtue of having the neurology we do.²⁴³ It does not mean that that concept is present from birth, but it also means that they are not learned as that would require acquiring them from nothing. For example, we are disposed to acquire the concept DOORKNOB on encountering doorknobs. Had we never seen doorknobs the concept would not be acquired. What matters here is that there is some form concept acquisition.

'5. Concepts are *public*: they're the sorts of things that people can, and do, *share*'.²⁴⁴ This allows for communication in so far as someone can cause another to entertain the same concept as them. It also means that people can entertain the same thought.

On this view concepts are physically realised and play the role of categories that can be employed to think about objects in the world. That they are physically realised lets them causally interact with the agents having those thoughts in ways that reflect the content of that thought. The character of a concept might be sufficient to count as its meaning and this plays a role in determining its content. Concepts refer to things in the world and different individuals can entertain the same concept. In the next section I will consider sentences in the language of thought.

2.5 A Sentence in Mentalese

Thoughts are sentences in Mentalese. A sentence in Mentalese is composed out of concepts. They are subject to computations in virtue of their logical form, and this allows for truth preserving inferences which preserve the semantic features of that sentence. These representations explain the

²⁴² Ibid 150

²⁴³ Ibid 146

²⁴⁴ Fodor 1998: 28

psychology and behaviour of the individuals that entertain them. It is in virtue of these mental representations that we can entertain propositions.

These Mentalese sentences are like natural language sentences in so far as they are composed out of parts and their composition determines, at least in part, the proposition that thought expresses. On this view, 'Mental states are relations between organisms and internal representations, and causally interrelated mental states succeed one another according to computational principles which apply formally *to the representations*'.²⁴⁵ So, to consciously entertain a given propositional attitude is to have a mental representation which relates the subject to the proposition that they desire/believe/hope for. Having a belief might be simpler, it may just be a case of having a mentalese sentence in one's "belief box".

One important feature which I will not be agreeing with is Fodor's rejection of intensions in favour of a purely extensional account.²⁴⁶ This is, in part, due to the definition of propositions as functions that was given in chapter one. A functional account of propositions is an intensional one, so Fodor's account would exclude this account of propositions. Following Fodor on this point would impede a discussion of context-sensitivity at the level of thought. Fodor's view does not allow for distinctions between various kinds of underdeterminacy because he denies that there is such a thing as content, in the sense of a function from possible worlds and indices to truth-values. (I will discuss some more substantial reasons against extensional accounts of content in chapter 5).

In the next section I will give an account of what I take to be Fodor's view in relation to the context-sensitivity of representations in natural language and mentalese. I will also contrast Fodor's view with others that are available in the literature, such as Radical Contextualism (RC, according to which all or nearly natural language sentences are context-sensitive) and RRC.

²⁴⁵ Fodor 1975: 198

²⁴⁶ Fodor and Pylyshyn 2015: 1

2.6 The Mixed View

When it comes to the relationship between thoughts and context-sensitivity, Fodor is often interpreted as a proponent of the Mixed View (MV).²⁴⁷ The MV makes two claims. The first is that all or almost natural language sentences are necessarily context-sensitive. This is claim (a) of RRC.²⁴⁸ The second is that mental representations, such as sentences in Mentalese, are not context-sensitive. This is to reject claim (b) of RRC, where (b) is the claim that all or almost all thoughts are context-sensitive. The semantics of thoughts does not (necessarily or otherwise) underdetermine the content (with a few possible exceptions that might be made for demonstratives). On Fodor's view, a thought just is its content and it seems that there is nothing else that determines content.²⁴⁹ The views on offer can be seen in the following table:

General View	Versions	All/almost all natural Language sentences are context-sensitive	All/almost all thoughts are context-sensitive
Mixed View (MV)	Fodor's LOTH, Radical Contextualists e.g. Carston.	✓	✗
Univocal context free	Semantic Minimalism or Moderate Contextualism	✗	✗
Univocal context dependent view	RRC e.g. Travis, Searle, Clapp, Jaque.	✓	✓
(Other) Inverted Mixed View	Unoccupied.	✗	✓

²⁴⁷ See Fodor (2001: 10-11), where he advocates the MV. Those who take him to be proponents of the MV include Jaque (2017) and Clapp (2013).

²⁴⁸ (a) all or almost natural language sentences are necessarily context-sensitive. (b) all or almost all thoughts are necessarily context-sensitive.

²⁴⁹ Fodor 2001: 14

The view is described as “mixed” because it mixes context-sensitive natural language expressions with context insensitive Mentalese sentences.²⁵⁰ In earlier work Fodor seems happy to allow for token-underdeterminacy at the level of thought in those cases in which natural language words are similarly subject to token underdeterminacy. He claims that ‘(the machine ought to be unable to determine whether the predicate applies) IFF (the speaker is unable to determine whether the predicate applies)’.²⁵¹ However, in later work Fodor seems to be more hostile to underdeterminacy in thought. I will elaborate on these quotes below. For the moment I will take Fodor to reject token underdeterminacy as a pervasive feature of thoughts. So, a Mentalese sentence will express a content and that will be a total function.

2.6.1 Comparisons with the Mixed View

This mention of the MV invites comparison with other kinds of views, two of which are worth mentioning at this point. The first would be a univocal non-context-sensitive view, according to which neither natural language sentences, nor thoughts, are *pervasively* context-sensitive. On this view, both (a) and (b) of RRC are false. It is not the case that all or almost all sentences in natural language are context-sensitive. It is also not the case that all or almost all thoughts are context-sensitive.

2.6.1.1 Semantic Minimalism

Spelled out in this way, it seems that one could be a Semantic Minimalist and still accept the view that neither thought nor language are pervasively context-sensitive. Semantic Minimalists deny (a) and it is compatible with their view to deny (b). Indeed, it is important that (b) would also be denied on Borg’s understanding of Semantic Minimalism, as her view makes use of modularity to explain how we are able to comprehend utterances automatically, seemingly without making use of

²⁵⁰ Jaque 2017: 4. Whilst Jaque confines her definition of the MV to type underdeterminacy, I do not wish to restrict the MV to this form of context-sensitivity only. Instead, I want to allow that token-underdeterminacy is also excluded on the MV.

²⁵¹ Fodor 1975: 63

additional information.²⁵² Getting to more complex cases of speaker meaning requires making use of non-modular inferences, that is, inferences that might, in principle, include any premise available to the interlocutor.

So, when a speaker describes a party by saying “There was not enough drink and everyone left” we need to make some inferences on the basis of what we know about parties and the speaker to understand them properly.²⁵³ Here, relevant information would be about the sort of drink that is typically drunk at parties, often alcoholic, whether the speaker would consider that good or bad, and why the speaker would mention that everyone left. Depending on the information, one might conclude that the speaker meant that there was not enough alcoholic drink, and everyone left early, possibly making it a bad party for the speaker. However, the Minimalist maintains that the original sentence conveys a proposition even independently of these additional inferences. It would follow that the thoughts we have that concern the literal meaning of an utterance in natural language (in a language we understand) would only be minimally context-sensitive. There may be other, more context-sensitive thoughts on Semantic Minimalism.

However, this spelling out of the theory seems to miss the point. The point of a univocal context insensitive view would be to deny that there is *any* context-sensitivity in either natural language or thoughts. But such a view seems too strong to be true. One need only point to cases of indexicals and demonstratives in natural language. So perhaps the best way to explain this view is to say that it only allows for a limited degree of context-sensitivity in natural language, and similarly limited context-sensitivity in thought. This is in keeping with what Semantic Minimalism suggests.

2.6.1.2 Really Radical Contextualism

A third view would be one that I have already mentioned, namely the RRC view. This is the view that all or almost all natural language sentences are context-sensitive, and the same is true of thoughts.

²⁵² Borg 2004 Ch. 2

²⁵³ Borg 2016: 339

Most criticisms of the MV are made from the perspective of RRCs, or at least people who maintain that if natural language sentences are pervasively context-sensitive, then so are thoughts.²⁵⁴

2.6.1.3 *Inverted Mixed View*

A final position, which is to my knowledge unoccupied, would be an inverted Mixed View. This view would claim that not all natural language sentences are context-sensitive, and the claim that all thoughts are context-sensitive ($\sim(a \wedge b)$). This view is also compatible with Semantic Minimalism in so far as both accept ($\sim a$). However, this inverted mixed view makes our comprehension of natural languages more complex. There is also an issue of how we would account for language being more explicit, or less underdetermined, than the thought we'd seek to express with it. It would make our understanding of our own natural languages even more mysterious, as the content of our thoughts have less information in them than the natural language sentence. That would seem to suggest that there are parts of our natural language sentences that we do not understand or would otherwise struggle to grasp.

2.6.2 *Mixed View*

In what follows I will elaborate on the MV and some of Fodor's arguments in favour of it. I will then discuss, in general terms, some of the criticisms that have been made of the MV, before mentioning some of the shortcomings in these arguments.

There are a few places in which Fodor seems to support the MV. One of the clearest is when he states that:

...the content of a [natural language] sentence may be inexplicit with respect to the content of the thought it expresses, a thought can't be inexplicit with respect to its own

²⁵⁴ So, claims are often made to the effect that 'If one endorses the arguments of radical pragmatics against the compositionality of language, then one should also reject the compositionality of thought' (Clapp 2010: 299), 'I will argue that their arguments fail to establish that representations that are free of Type-Underdeterminacy are indispensable. As a consequence, alternative approaches are, at least, tenable' (Jaque 2017: 4).

content; there can't be more—or less—to a thought than there is to its content because a thought just *is* its content.²⁵⁵

Here Fodor is endorsing the second claim of the MV, whilst allowing the possibility of the first by claiming that natural language is inexplicit. The idea would be that any inexplicitness regarding content in natural language can be removed by using the context to make an inference to an explicit thought.

Fodor makes a similar claim elsewhere. He also claims that 'Statements express thoughts, and the content of thoughts isn't constituted by their contexts... All context can do is provide the hearer with more or less reliable information about what thought that was'.²⁵⁶ Here, Fodor still accepts the second tenant of the MV, but it is not clear that he accepts the first one, that natural language is context-sensitive. Rather, he accepts that context may be used to infer the intended speaker meaning, rather than claiming that context is necessary for that sentence to express a proposition at all. When it comes to the explicitness of thoughts this rules out the possibility that thoughts can benefit from the context.

This suggests that thoughts may not admit of any form of underdeterminacy. Either the context cannot help us to resolve partial functions or it can. It seems that there are cases in which we can make use of the context, possibly in the form of more information about an object, to resolve a partial function. For instance, we might learn that a feline creature that spoke Latin is descended from other normal cats.²⁵⁷ This might sway us to judge that it is a sort of cat. Fodor's denial of context-sensitivity at the level of thought might also lead him to reject that the context can play this sort of a role. In which case, it appears he must reject that the thought was ever partially underdetermined in the first place, or that the input was the issue. There is a risk of saddling Fodor

²⁵⁵ Fodor 2001: 14

²⁵⁶ Fodor 2005: 108

²⁵⁷ Recanati 2004: 142

with an unattractive view that is not his own, in which case one can say that this MV is a Fodor like view, which serves as a useful contrast to the more context-sensitive views that I wish to discuss.

However, given that Fodor thinks that it is thoughts that have content in the first instance, not natural language sentences, it seems reasonable to attribute the first part of the MV to him as well. Fodor has claimed that ‘as a matter of empirical fact, language is pretty clearly *not* compositional; so it can’t have content in the first instance’.²⁵⁸ This leads Fodor to claim that it is thoughts that have content in the first instance.²⁵⁹ In more recent work he also says that ‘forms of speech inherit their semantic properties from those of the thoughts they are used to express’.²⁶⁰ If natural language is not compositional then its parts don’t determine its meaning. On the picture Fodor paints here, natural language might just be a tool to allow a hearer to work out what the speaker intends to communicate and (a) might be true. This intended content might be a context insensitive thought (or thoughts).

The MV follows from some of the features of the LOTH that I have mentioned. These features include the composition of thoughts; there is also a regress argument which I will discuss below. Finding a space that can accommodate both the LOTH and context-sensitivity would be good evidence in favour of the thesis. If context-sensitivity of thought can avoid these criticisms whilst not giving up too many of the tenants of the LOTH, then I will take this to be a good reason to believe that there is such a logical space.

2.6.2.1 Debate Around the MV

The debates about the MV tend to focus on the compositionality principle, understood as the claim that the content of a representation depends only the parts of that representation and the way in which they are combined, and Fodor’s claim that thoughts must be fully explicit regarding their content.²⁶¹ If the criticisms of these claims are successful then it is possible that thoughts are context-

²⁵⁸ Fodor 2001: 11

²⁵⁹ Ibid 14

²⁶⁰ Fodor and Pylyshyn 2015: 12

²⁶¹ See Jaque 2017, Elguardo 2005, Clapp 2013

sensitive. It does not prove that they are yet. Though some claim that, in the absence of a barrier to thinking of thoughts as context-sensitive, and believing natural language to be pervasively context-sensitive, we should think of thoughts as context-sensitive, but only in the absence of arguments for the MV.²⁶² Even if this is correct, however, it does not tell us how we can understand the context-sensitivity of these thoughts. What they do is to disarm some of the objections to a RRC view and allow for the possibility of context-sensitivity at the level of thought. They do not, however, describe the way in which thoughts should be understood as context-sensitive.

In what follows, I will outline some of the objections that have been made to Fodor's version of the MV. These arguments focus on undermining Fodor's arguments for the MV. Fodor's aim to show that Really Radical Contextualism is not possible, so these arguments can also be taken as a defence of the MV. I will use this to show some conclusions that have been drawn in the debate so far.

Importantly, whilst many of these points appear successful, they do not respond to all the arguments that Fodor makes in sufficient detail. I will then introduce the regress argument and it will be a central concern of this thesis to show that this regress can be responded to. In rough outline, the following objections have been made by opponents of the MV. This is not an exhaustive list.

2.6.2.1.1 Clapp

Clapp focuses his discussion of the MV on the premise that Fodor uses to support it. This is the premise that 'As between language and thought, *at least one of them* is compositional'.²⁶³ Here, compositionality means that the content of that representation will be determined by the syntax and content of the parts and nothing else. So, whichever is compositional is not context-sensitive.

Successfully targeting this premise does serve to undermine an argument in favour of the MV. The argument is, roughly, that either thoughts are compositional, or natural language sentences are. At

²⁶² Clapp 2012: 320

²⁶³ Clapp 2012: 300

least one of them is compositional. Natural language sentences are not compositional. Therefore, thoughts are compositional.²⁶⁴

Clapp's counter arguments, if successful, will undermine this argument for the MV. Clapp makes several counter arguments to undermine the claim that at least one of either natural language sentences or mentalese sentences is compositional. To do so, Clapp aims to show that Fodor's argument that thoughts cannot fail to be compositional is invalid. The argument given by Fodor is that a thought is identical with its content. If that is correct, then no two tokens of a given mentalese sentence *m* can have different contents. In which case, it seems that thoughts are compositional.

Clapp's counter argument is that thoughts cannot be understood as both content and representation.²⁶⁵ This is because "thought" can be understood in two ways. On one reading, "thought" does just mean content. On another, "thought" is the vehicle of content. On the first reading Fodor's claim is trivially true. On the second there is a more substantial claim, that the vehicle of content is also that content. However, for Fodor to make the claim that the vehicle, or the representation, is its content is question begging in this context.²⁶⁶ RRC's argue that there is at least one kind of underdeterminacy that applies to thoughts. If thoughts are type- underdetermined then they are not identical with their content, for instance. Fodor's claim that a vehicle is its content is not a premise any contextualist would have to accept without independent argument. In this case Fodor has not proven that mentalese sentences are identical with their content, so there is not a reason to accept that thoughts need to be any more compositional than natural language.

Whilst Clapp's argument is interesting it is not sufficient to establish a non-Mixed View. There are other arguments in favour of a MV (see below). Also, whilst Clapp's argument undermines a few reasons for thought being compositional that does not prove that thought is not compositional. Though Clapp does make a case for the non-MV as well, it is dependent on what I will call "the

²⁶⁴ Fodor 2001

²⁶⁵ Clapp 2012: 306

²⁶⁶ Ibid 306-7

concerning relation". This relation will be most relevant when discussing Searle's view in chapter 4. I will discuss this issue in more detail then.²⁶⁷ That discussion will differ from Clapp's as it goes into more detail on this relation to try to show that it is an alternative that proponents of the MV should take seriously.

2.6.2.1.2 Jaque

Jaque argues that proponents of the MV '...fail to establish that representations that are free of type underdeterminacy are indispensable'.²⁶⁸ Here Jaque is targeting an argument from creativity that Fodor makes. The idea is that human thought is creative, in that we can, in principle, produce an unbounded number of thoughts, and these thoughts can be entirely novel to us. Fodor argues that to be able to produce thoughts like this, thoughts must be productive in that their components can be combined in systematic ways to produce unboundedly many thoughts.²⁶⁹ For instance, the thought THAT IS A NICE VIEW can be made into the thought I THINK THAT IS A NICE VIEW. From here, it is easy to see how this process might lead to an unbounded number of thoughts. For this to work, Fodor argues that a thought's content must be composable from the content of its parts and the manner in which those parts are combined.²⁷⁰ (Natural language does not compose for Fodor and he assumes that one or the other will have to be composable, as there is not an additional source of content and he does not think it is plausible that neither has content).²⁷¹

Jaque argues that composition of content is not the only way to account for productivity of thought, however. She claims that '...if it is creativity we are interested in, or our ability to think new thoughts, then the best explanation given the scenarios of underdeterminacy involves the creation of *ad hoc* concepts, as Carston defends. But then productivity is not the key to our ability to think new thoughts—the creation of new concepts is'.²⁷² Here Jaque is denying that productivity is required to

²⁶⁷ See section 4.3

²⁶⁸ Jaque 2017: 4

²⁶⁹ Fodor 2001: 6-7

²⁷⁰ Ibid 14-15

²⁷¹ Ibid 2

²⁷² Jaque 2017: 8

explain our ability to form new concepts. If productivity is not required, then compositionality of content is also not required. This is the form of Jaque's argument.

An *ad hoc* concept is a concept that is created by altering a previously existing one, by a process of modulation.²⁷³ This process occurs when we find that we need a new version of a concept, either to make sense of what someone has said or to navigate the world. An example would be the change of SAINT to SAINT*. Here the concept of SAINT is altered to give a new concept with a different content, in this case someone who is kind rather than someone who has been canonised. Using the concept SAINT* in place of SAINT seems to be a case of thinking a novel thought. Hence, we can be creative in our thinking without having to be productive.

There are some potential shortcomings with this argument. One is that it does not show that we can use *ad hoc* concepts in all cases of creative thinking. It may be that there are cases in which we would want to think of a brown cow, having never done so before. Here the explanation that one takes the concept BROWN and conjoins it with the concept COW seems to be at least as plausible as the explanation that we form a new concept COW* which applies to only cows that are brown. In which case it seems that there is still a good case for thinking that productivity plays an important role in our ability to think creative thoughts.

In response Jaque maintains that 'Given Fodor's notion of productivity, meaning compositionality is sufficient for a system of representations to be productive'.²⁷⁴ Here meaning compositionality is the claim that 'The meaning of a well-formed declarative sentence S is determined by the meaning of the expressions in S and the syntactic structure of S'.²⁷⁵ So, even if the content of a thought is not something that can be determined by the thought alone, the thought's meaning can be. Importantly for Jaque's point, meanings can be combined in systematic ways to account for productivity that is not a result of the compositionality of content. For Fodor, this should also be sufficient for

²⁷³ See Wilson and Carston 2007, Allott and Textor (2012).

²⁷⁴ Jaque 2017: 8

²⁷⁵ Ibid 7

productivity in English as he claims that English is productive whilst it does not have content.²⁷⁶

Having meaning compositionality without content compositionality suggests that thoughts can be type – underdetermined, as the content is not determined by the thought, that thought’s meaning is compatible with several different contents.

If Jaque is correct that we do not need thoughts that are free of type-underdeterminacy then the way is more open for a non-MV. This is because it removes an important objection to a RRC view. It does not establish that a RRC view is correct or independently plausible. Though she has also argued that mental representations can be subject to token-underdeterminacy, meaning that their content does not determine a truth-value for some inputs, she does so in a way that is very similar to Travis.²⁷⁷ I will discuss token-underdeterminacy in more detail in chapter 5.

2.6.2.1.3 Elguardo

Elguardo seeks to undermine Fodor’s case for the explicitness of thoughts, on the basis that

‘...compositionality does not require that complex symbols be explicit about their semantic contents’.²⁷⁸ Elguardo focuses on the case of English as a language that has compositionality whilst it

is not explicit. To be explicit, a representation must meet (at least) two constraints. **‘Constraint [1].**

*Each syntactic constituent of a meaningful (non-idiomatic) complex symbol corresponds to a constituent of the symbol’s semantic content. **Constraint [2].** Each constituent of the semantic*

*content of a meaningful (non-idiomatic) complex symbol corresponds to some syntactic constituent of the symbol’.*²⁷⁹ If Elguardo is right that English can be both compositional and inexplicit, then it is not

the case that all language-like representations will need to be explicit in their content. In which case, it is possible that thoughts are language-like and inexplicit. They might be underdetermined in an interesting way.

²⁷⁶ See Fodor 2001: 7 and Fodor 2001:11-12, respectively.

²⁷⁷ Jaque 2017: 20

²⁷⁸ Elguardo 2005: 60

²⁷⁹ Ibid 65

Fodor argues that if English is compositional then sentences of English will be explicit in the content that they express.²⁸⁰ For Fodor, ‘...the content of a sentence is, plus or minus a bit, the thought that it is used to express’.²⁸¹ So, Fodor can be understood as claiming that English sentences are not explicit with regard to the thought that they express. Therefore, English is not compositional.

Elguardo’s response is to deny the conditional claim, which he labels [P1], that ‘If English has a compositional semantics, then *Constraints [1] and [2]* hold for each complex (non-idiomatic) English symbol that has semantic content’.²⁸² Elguardo considers a few examples of English to make this case. One is the sentence “Some boy fell”. Elguardo argues that the “some” in this sentence ‘...doesn’t correspond *on its own* to any semantic constituent relative to a semantic interpretation’.²⁸³ However, the syntactic form “Some boy fell” can be ‘...paired with a compositionally determined meaning given the meanings of its immediate constituents and its syntactic structure’.²⁸⁴ In which case, the sentence is compositional whilst failing to meet constraint 1, so the conditional [P1] is false. Elguardo also considers other cases of English sentences that appear compositional yet fail to meet the necessary conditions for explicitness. He also considers ‘...expletive uses of “it”, as in “It is true that Bush was elected to office in 2000”, and “there”, as in “There appears to be a cat hiding behind the couch”, which are syncategorematic since they don’t designate anything from any semantic category. In both cases, nothing in the content of the thoughts expressed by those sentences (relative to a context) corresponds to “it” and “there”, respectively’.²⁸⁵ Other examples are prepositions and words that undergo type-shifting.²⁸⁶

If Elguardo is right that representations can be compositional (in the sense that their meaning composes) whilst being inexplicit then it is not true that representations will have to be explicit if

²⁸⁰ Fodor 2001: 12

²⁸¹ Ibid 11

²⁸² Elguardo 2005: 66

²⁸³ Ibid 68

²⁸⁴ Ibid 68

²⁸⁵ Ibid 71

²⁸⁶ Ibid 71, 69

they are compositional. This is important as it seems that meaning compositionality is needed to account for the productivity of thought in at least some cases. This opens the way for thoughts to be inexplicit about certain features. This inexplicitness may open the way to underdeterminacy at the level of thought. If constraint 2 is not met by thoughts, then it seems that they have more semantic content than they do syntactic parts. The idea would be the thought's content is made explicit by the context. That not all representations need to be explicit if they are to be compositional does not, of course, do much to show that thoughts are like this. It does, however, show that such a view is possible.

2.6.3 Summary

A common theme in these papers is that they do a lot to remove obstacles to an RRC view. However, I would suggest that there are still some arguments against RRC that have not been addressed in detail. For instance, there is a worry that context-sensitive thoughts will lead to a regress. I will discuss this argument below. These papers also do not go into detail about what a context-sensitive account of thoughts might look like. Though, they do gesture to certain accounts of thought (for instance, Clapp discusses a concerning relation and discusses Travis cases).²⁸⁷ One of the main aims of this thesis will be to spell out what some of these views might look like in more detail.

On Fodor's view, such context-sensitivity is absurd. When considering an utterance of "I shot an elephant in my pyjamas" Fodor argues that a context-sensitive view of this thought '...conjures up a situation more absurd than an elephant in pyjamas. "I wonder what I meant by saying that? I shall inquire into the context of my utterance in order to find out. Since, patently, no such situation can arise, it would seem to follow that what Groucho meant cannot be constituted by any contextual fact'.²⁸⁸ The criticism here will take some unpacking to get right, but for the moment we can sketch Fodor's objection as saying that if context-sensitivity exists at the level of thought, we are left in a

²⁸⁷ Clapp 2013: 302, 320-321

²⁸⁸ Fodor 2005: 107-8

situation in which people have to work out what they have thought. Fodor takes this to be absurd and so objects to the plausibility of RRC.

2.7 The Thesis Revisited

We can now see a little more clearly what the driving question that I will be addressing is. The question is whether some context-sensitivity can be accepted in the language of thought, even when these language of thought sentences are foundational. This claim is summarised as:

Weak foundational context-sensitivity (modal): Some foundational sentences in a language of thought could be context-sensitive.

On this view it is possible that there is a foundational sentence in Mentalese which can express different propositions, or a different set of truth conditions (have a different content) depending on the context in which it is thought. This means that it would be possible for there to be the same concepts in the same ordered structure (the same Mentalese sentence) and for it to produce a different function when tokened in different contexts. This will be true of at least some Mentalese sentences. If this is not possible, then the hypothesis is false. A test of whether the function will have changed will be to see whether the truth-value changes when the thought has the same input.

I will not be considering all kinds of thoughts as potential candidates for proving this thesis. For instance, thoughts that are not truth-evaluable will not be considered. So, when we ask ourselves questions, or tell ourselves what to do, there cannot be a truth-value to these thoughts. They will not be considered. This is not to say that there are not interesting questions to ask here. For instance, the thought GET OUT OF BED might on some cases be an imperative directed at oneself, or a thought of frustration directed at a lazy friend. I will try to focus on thoughts that can be taken to straightforwardly express a proposition. Thoughts that are triggered by someone else's assertions are a good example of this, but they needn't be the only category.

2.7.1 *General Difficulties with Context-Sensitive Thoughts*

There are several concerns that will recur in this thesis. One difficulty is to show that a given thought is foundational. So, if we have a demonstrative thought THAT SHADE OF RED IS QUITE STRIKING, the task would be to show that the truth-value of this sentence is not dependent on some other thought like THE SHADE OF RED PRODUCED AS A RESULT OF LIGHT OF WAVELENGTH X STRIKING NEIL'S EYES IN NORMAL CIRCUMSTANCES IS QUITE STRIKING. If the first, demonstrative, thought, were dependent on the second thought, then it would seem to lose whatever context-sensitivity it had in the first place. If the first is context-sensitive but the second is not, and the first is dependent on the second, then there is not context-sensitivity at the foundational level.

This is not, however, a problem for Really Radical Contextualists. At least, not if they are correct that all representations, and therefore all sentences in Mentalese, are context-sensitive. This would mean that even a foundational thought would still be context-sensitive. RRC arguments often attempt to show that the representations that humans use cannot be context insensitive. So, the problem is mute for RRCs, if they are right. Proving this correctness is, however, its own problem as the possibility of a RRC view is still in question.

2.7.2 *The Regress Argument*

A second problem, which seems to occur when we have a context-sensitive foundational thought, is that it leads to an infinite regress. Carston claims that there cannot be a pragmatics of thought as 'thoughts are, in effect, an end-product' of pragmatic processing and if pragmatics does not end there then it will not stop.²⁸⁹ This argument is also hinted at by Fodor who says that 'All we care about is whether equivocation goes on forever; i.e. whether there are any *unequivocal* expressions of finite length'.²⁹⁰ The implication here being that if thoughts are context-sensitive then any one

²⁸⁹ Carston 2002: 76

²⁹⁰ Fodor 2004: 111 ft20. Jaque (2017: 87) seems to have a similar point in mind. However, she presents the argument as follows: two uses of an equivocal sentence can and often do differ in truth conditions. The only thing that can resolve the equivocation (=account for the difference in truth-conditions) is a mental representation-type which conceived as non-equivocal. Therefore, there must be non-equivocal mental representations. She then argues that equivocal mental representation-types can be used to resolve equivocation. Whilst I will argue that this is true, I will do so by considering theories of context-sensitivity more generally than Jaques, who focuses on specific examples. I have also put more emphasis on the way that

thought that we have might be equivocal between two contents. In which case, there will never be a non-equivocal thought. However, it seems that whenever one has such an equivocal thought, one should try to make it non-equivocal. Therefore, getting a non-equivocal thought is an unending process. One way of putting this argument goes as follows:

1. Foundational thoughts are not true in virtue of the truth of any other thought. (If any thought M is true in virtue of the truth of another thought Mn, then M is not foundational).
2. If a thought (=tokened sentence in Mentalese) M1 is context-sensitive, then there must be a process by which the context affects the thought.
3. This process of removing context-sensitivity will result in a new thought, M2, which takes the context into account.
4. M1 is true in virtue of the truth of M2 (in that context).
5. All thoughts are context-sensitive (RRC).
6. Therefore, M1 was not foundational. (1, 4)
7. M2 was not foundational (1, 3, 5).

Here (7) follows from the premises because if (5) is true, M2 is also a context-sensitive thought. In which case there needs to be a way of resolving the underdeterminacy that thought has. If (3) is true, and thoughts are the bedrock of representation, then M2 can only result in a new thought. That thought will also be underdetermined, if (5) is true and the process will continue *ad infinitum*.

This presents a challenge to my core claim that there can be context-sensitive thoughts. If no context-sensitive thoughts can be foundational, then there could not be any foundational context-

equivocal thoughts lead to a regress as I think that this better reflects concerns that are had by proponents of the MV, see Carston (2002: 76) who claims that 'Having thoughts is a strikingly different kind of mental activity from comprehending utterances. We do not have to undergo a process of *comprehending* occurrences (tokenings) of our thoughts, as we must comprehend occurrences of linguistic utterances. Thoughts are, in effect, an end-product...'. She goes on to say that if this were not so then we would have a regress (Carston 2002: 76).

sensitive thoughts. This is a problem, as it means that there are no thoughts that are true in virtue of themselves. Instead, all thoughts are true in virtue of the truth of some other thought. As each new thought attempts to remove the underdeterminacy that was in the previous one, we never get a context insensitive thought. Instead, each new thought needs some other thought and it is true in virtue of the second thought.

One alternative way of spelling this concern out might go as follows. We want to know why it is that a given thought (M) represents the state of affairs that it does. One response might be that it is interestingly connected to some other thought (Mn), in that Mn is the foundation of M. But that only moves the question. Now we want to know why Mn can represent a given state of affairs. We don't want it to be the case that Mn stands in relation to some other thought, as then all we do is to move the question again. So, we may want to say that Mn (or Mn+x) is foundational. But this seems to be precisely what RRC is saying we cannot do. Mn cannot just represent that state of affairs if it is underdetermined, as it will be able to express other functions which are incompatible with that state of affairs being true.²⁹¹ Once we have reached thoughts then it seems that we cannot appeal to a further meta-language. So, we must go from thought to thought without any ever being foundational, that is be able to be true of a given state of affairs without appeal to another thought. In this case, there is a regress of equivocal thoughts.

If thoughts are all token underdetermined, then there is a similar problem. It seems that there are two possible extensions for a given input, true or false. Arguably, when one picks which of these is appropriate, one is moving from an underdetermined thought to one which is not underdetermined for this input. Yet the next thought will also be underdetermined (by hypothesis). So, there will be cases in which it is not foundational either. For example, we may have a thought:

(m) X IS A CAT.

²⁹¹ This seems to be the point that Travis (1997, 2006) makes.

Where X appears to be a typical feline. We then discover that X speaks Latin.²⁹² We may decide that our initial thought is still true of X, though it seems initially underdetermined. Now that it is no longer underdetermined, we arguably have replaced the thought with a new one:

(Mn) X IS A CAT*.

Here, Mn seems to be the foundation of M in that context. However, we may then discover that not only does X speak Latin, but X also radiates some hitherto unknown form of radiation. Now we are not sure whether CAT* applies. It seems that we might then to develop the concept into CAT**, and so on, for several different possible worlds. The cat may also spontaneously combust, resurrect nine times, or be a natural when it comes to deductive inference. For each new input we may need a new concept to determine whether X, or any other cat like animal, is still a cat when we realise it has these properties. In this case, no thought is foundational across all possible inputs, but only for certain inputs. Or it may be that we cannot give a value to these thoughts, in which case they remain indeterminate and have no truth-value. In which case, we do not have a foundational thought as there is not a true thought for that input.²⁹³

This argument could also be taken to support an explicitness claim about thoughts, along the lines that Fodor suggests. The idea would be that each thought will need to be explicit about its own content, or it will lead to a regress. There are various responses to this argument. One obvious response is to deny either (1) or (5), as both are needed to generate the regress. However, (1) seems to be integral to what it is for a thought to be foundational. So that is difficult to give up on. If one gives up on (5), there still needs to be a way of maintaining the thesis (I will consider character-underdeterminacy in chapter 3). This will still result in having to have other thoughts that are

²⁹² Recanati 2004: 142

²⁹³ This may be a bullet that some would be happy to bite, but it seems rather unsatisfying to simply shrug our shoulders in such cases. It seems that there is a real question as to whether X is a cat when it speaks Latin, and that we should have some sort of an answer to this.

foundational and are context-sensitive. So, whilst one could give up on (5) that would not solve this problem, certainly not for RRC.

This leaves premises (2), (3) and (4) as possible candidates for rejection. It seems that RRCs typically reject (3). The challenge will then be to make sense of context-sensitivity which does not result in a new thought. So, thoughts, alone, cannot be the bedrock of representation on a RRC view. The issue then is to say what is. This will be a general theme throughout the thesis.

An alternative response to this line of argument would be to deny that there is any hierarchy of thoughts. In this case, there are no foundational thoughts. So, it is not a problem that we are unable to achieve them. This approach would also remove some of the obstacles to establishing context-sensitivity at that level of thought. For instance, when it comes to indexical and demonstrative thoughts one need only establish that these thoughts are possible, and not that they do not depend on some other saturated thought.

However, it is not obvious that this is the best approach. For one, it might not be sufficient to avoid a regress. It would just avoid this version of the regress. Even if what we are looking for is not a foundational thought, there is still a challenge to make sense of how the context affects either the content or the extension of a thought without generating a new thought. Even if a foundational thought is not what one requires, it seems that this remains a problem for the RRC.

A second reason for not taking this approach is that it risks misconstruing the nature of context-sensitivity as something that is inessential to thoughts. If there are no foundational thoughts it might still be that thoughts which are context-sensitive are merely there as a convenience. The idea may be that we off load some of the work onto the world, whilst we remain fully capable of entertaining context insensitive thoughts. This would misconstrue the RRC position, according to which context-

sensitivity is an essential feature of thoughts. If context-sensitivity is merely convenient, rather than essential, it is a much less interesting phenomenon.²⁹⁴

2.8 Summary

To summarise this chapter, I have given an exegesis of the Language of Thought Hypothesis (LOTH). This is the thesis that thoughts can be understood as language like in so far as they have a syntax. The LOTH also implies the representational theory of mind (RTM) and the computational theory of mind (CTM). According to RTM, we token mental representations and these representations express propositions. This is how we entertain propositions. These mental representations will be in Mentalese. According to CTM, thoughts can be identified and operated on purely in virtue of their syntax. The scope of computational processes is limited to mental modules. A general processor, whilst it makes use of the language of thought, does not operate computationally. This is to avoid having to encounter the Frame Problem.

So, we have a view of thoughts according to which thoughts are language like, they are physically realised in the brain (by neural activity), these representations will express propositions, and the representations have syntactic properties that allow them to be operated on in a computational manner.

I have also given an account of the MV, according to which thoughts cannot be context-sensitive. Yet natural language sentences are pervasively context-sensitive. I attribute this view to Fodor. The arguments in favour of this view have been criticised by several opponents. All this criticism goes to support the RRC's view, that both natural language sentences and thoughts are context-sensitive.

However, removing some of the obstacles to the RRC view is not enough to establish that it is true. Nor does it tell us which version of RRC would be the preferable one. There is a further problem in that there is the threat of a regress for RRC which is one fundamental obstacle that has not been

²⁹⁴ A version of the distinction between context-sensitivity in natural language that is merely convenient and that which is essential can be found in Carston 2002: 29

dealt with. This is the regress argument. This is the primary issue that I will aim to address. A second issue is whether (b), that all (or almost all thoughts) are necessarily context-sensitive, is compatible with a CTM.

3. De Se Attitudes and Computation

Essential indexicality in natural language is the ‘dominant’ position in the literature.²⁹⁵ Essential indexicality is the claim that certain indexical attitudes cannot be replaced by non-indexical attitudes without losing an important feature of that attitude.²⁹⁶ A second position is *de se essentialism*, which states that certain attitudes that one has about oneself (the *de se attitudes*) cannot be replaced by attitudes of a different kind without a loss or change in what that attitude is able to explain. This view is often defended on the basis that *de se attitudes* pose a unique problem for theories of propositions.

De se attitudes are often taken to have their own kind of content. *De se contents* are contents that make a perspectival or first-personal reference to the person entertaining that content and are usually expressed in English with “I”. *De se attitudes* are often taken to need an indexical element. If *de se attitudes* do have an indexical element, then it seems that *de se thoughts* (thoughts tokened when having a *de se attitude*) will also be essentially indexical. When one has a *de se attitude* one has a *de se thought*, understood as a sentence in mentalese e.g. I AM IN READING, and the thoughts have their own kind of content which makes a reference to the individual thinking it.

I will focus on *de se essentialism* as including an indexical element because it is promising for the essential indexicalist and has come under pressure in recent debates.²⁹⁷ *De se essentialism* is a view that has been targeted by *de se sceptics* who deny that there is anything essential or special about *de se attitudes*. *De se essentialism* can provide a useful test case for the possibility of indexicals in thought, provided that *de se attitudes* are also ones that include an essential indexical element.

That thoughts include an indexical element is a view I will call “indexicalism”. I will not be concerned with indexicalism more generally, though if indexicality is possible in the case of *de se attitudes* then this possibility supports the claim that indexicality is present in other non-*de se attitudes*. I am not

²⁹⁵ Cappelen and Dever 2013: 2

²⁹⁶ Recanati 2012: 33

²⁹⁷ See Ninan 2015, Cappelen and Dever 2013, Lewis 1979, Perry 1979, Millikan 1990, Bermudez 2017.

looking at demonstratives partly due to considerations of space, and partly because demonstratives and *indexicals* face the same problem in the CTM, in that there will be different tokens of a given type, that nonetheless have different contents. I will elaborate on that problem in more detail below.

Here I want to discuss three widely held views in the literature. These are, first, *de se essentialism*, understood as the claim that some thoughts that we have about ourselves cannot be replaced with an alternative thought without some important loss. Secondly, the view that these *de se thoughts* require an indexical component. I will call the conjunction of these views *de se indexicalism*. The third is the computational theory of mind (CTM). It has been pointed out that there is a tension between these views, as computations struggle to treat different tokens of the same type of symbol differently, but this treatment is what indexicals require.²⁹⁸ Whilst this seems to make indexicals incompatible with the CTM, I will argue that *de se indexicalism* is compatible with the CTM. This is significant as it opens the door to an indexical at the level of thought, “I”. This suggests that there can be indexical thoughts and therefore context-sensitivity at the level of thought. In sections 3.2-3.4 I will present and defend each claim.

In section 3.5 I argue that the apparent tension between *de se indexicalism* and the CTM can be avoided. I claim that the alleged inconsistency depends on a certain understanding of the CTM according to which computation is the only means of processing information and reasoning. I will follow Fodor (2000) in denying that we should accept such a strong view of the CTM. Instead, the CTM need only commit one to the claim that *some* processes and reasoning undergone by the mind are computational. On this weaker understanding of the CTM we do not need to maintain that all the processes in the mind are computational. This opens space for non-computational processes in the mind. Even if we allow that computations cannot deal with indexicality, we do not require

²⁹⁸ Ball 2010. When it comes to homonyms and proper names, it may be that the best way to understand these on a computational theory of mind is to give each one an identifying subscript. So “bank” might become BANK1 and BANK2, for the financial institution and the riverside, respectively.

computation to do so to have *de se attitudes* that are indexical. I will conclude that this gives us reason to accept that context-sensitivity is possible at the level of thought.

3.1 Inconsistent Triad

In this section I will present three claims associated with *de se indexicalism* and the CTM. I will then discuss each claim and present an argument that suggests that they are inconsistent with each other.

The three, widely held, claims are:

- i) Having a *de se thought* is essential for representing ourselves self-consciously.
- ii) *De se* thoughts are indexical, in the sense that they have a context-sensitive (Kaplanian) character.²⁹⁹
- iii) Computations are only sensitive to the syntactic *type* of a symbol; they are not sensitive to other properties that the tokens that they operate on may have (Ball 2010).

3.2 Claim (i)

Claim (i) stems from work by Lewis (1979) and Perry (1977, 1979) in which they claim that some propositions cannot be grasped in a self-conscious way in the absence of a *de se attitude*.³⁰⁰ I use the term *de se attitude* to denote an attitude that involves a *de se thought*, that is a thought which makes reference to oneself in a significant way (to be spelled out below).

A *de se thought* is a thought that one has about oneself, but this condition is not sufficient. One also must realise that the thought is about oneself. Such thoughts are often expressed (in English) using “I”, though not always.³⁰¹

Evans adds two necessary conditions for thoughts to meet for them to be ones that we have about ourselves. The first is that one must recognise their significance to the ‘various special ways he has

²⁹⁹ Kaplan 1989: 506

³⁰⁰ Catañeda (1968: 446 -447) also discusses similar cases which motivate (i).

³⁰¹ Bermudez 2017: 12

(as every person does) of gaining knowledge about himself'.³⁰² So if one entertains a thought that is about oneself but does not appreciate that it can contribute to knowledge of themselves, then they arguably do not realise that that thought is about them. If Caesar had the thought that "Caesar is wanted dead by Brutus" and does not realise that he can know something about himself, he does not make the connection between "I am wanted dead" and "Caesar is wanted dead" to realise that he is wanted dead, then he has not realised that the thought is about him. So, it is not a *de se attitude*. A second necessary condition is that one '...must realise how to *act upon* propositions'.³⁰³ If Caesar thinks "Caesar will be assassinated" and does not realise how this belief should affect his behaviour, either to save himself or face his fate, then he does not seem to have realised that the thought is about himself.³⁰⁴ So, it is not a *de se attitude*.

Claim (i) is the claim that having a *de se thought* is essential for representing ourselves self-consciously, any other kind thought will not give the right kind of access to propositions.³⁰⁵ So the *de se thoughts* are essential to grasping these propositions. Any non-*de se thought* will not be capable of grasping that proposition in the correct way. Evans put the point as follows: 'Our self-conscious thoughts about ourselves are irreducible to any other mode of thought; in particular, they cannot be regarded as involving the identification of any object by any description'.³⁰⁶ So claim (i) is the claim that *de se thoughts* cannot be replaced by non-*de se thoughts* without losing the sense that the thought refers to oneself. In the following paragraphs I will present some of the original arguments for (i).

Some of Lewis and Perry's arguments are to the effect that we cannot explain some behaviour without positing a *de se attitude* (in which one entertains a *de se thought*). This relates to the second

³⁰² Evans 1982: 206, Recanati (2012/2017) makes a similar point about the SELF file depending on various epistemically rewarding relations that we have to ourselves.

³⁰³ Evans 1982: 206

³⁰⁴ If Caesar believes he is Brutus, then he has not realised how this should affect his behaviour.

³⁰⁵ For Perry (1979: 375) anyone can access a given proposition but not in any way. In this respect there is something that is essential (=irreplaceable) to the *de se thought*, it is just that the essential nature of these thoughts consists in the way we can access certain propositions, rather than which proposition we can access.

³⁰⁶ Evans 1982: 255

necessary condition that Evans mentions. So, the argument's form is that if two thoughts express the same proposition then they should have the same effect on the person's behaviour (all else being equal). When we have a *de se thought*, other *non-de se thoughts* do not have the same effect on the person's behaviour (all else being equal). Therefore, these two thoughts are not equivalent even though they express the same proposition. This suggests that there is something else unique to *de se thoughts*.

To motivate *de se essentialism*, Perry introduces a case in which he is following a trail of sugar in a supermarket, trying to tell the unfortunate shopper that their sugar bag is leaking. As he searches for the shopper:

...the trail [of sugar] became thicker. But I seemed unable to catch up. Finally it dawned on me. I was the shopper I was trying to catch. I believed at the outset that the shopper with a torn sack was making a mess. And I was right. But I didn't believe that I was making a mess. That seems to be something I came to believe. And when I came to believe that, I stopped following the trail around the counter, and rearranged the torn sack in my cart. My change in beliefs seems to explain my change in behaviour.³⁰⁷

Perry then wants to suggest that there is something special about this change in belief. Without the use of "I" in characterising the belief, it seems that we would no longer have an explanation of Perry's behaviour:

When we replace [I] with other designations of me, we no longer have an explanation of my behaviour and so, it seems, no longer an attribution of the same belief. It seems to be an *essential* indexical.³⁰⁸

Here Perry is claiming that when we seek to replace an "I" in certain thoughts, we will change the thought in a significant way. This is because we can only explain some changes in behaviour by

³⁰⁷ Perry 1979/2010: 366

³⁰⁸ Ibid 366

shifting from a non-*de se* expression (such as “a shopper”) to a *de se* expression (“I”). When Perry thinks of the person making a mess as “a shopper” he continues to search for them. It is only when he realises “I am making a mess” that his behaviour changes. It seems that even if Perry believed “Perry is making a mess” he would still need to have some sort of belief to the effect that “I am Perry” for that belief to make a difference to his behaviour.³⁰⁹

Perry supports the claim that we require some belief of the form “I am Perry” by using a different example. This is the case of Lingens the amnesiac.

An amnesiac, Rudolf Lingens, is lost in the Stanford library. He reads a number of things in the library, including a biography of himself, and a detailed account of the library in which he is lost. He believes any Fregean thought you think might help him. He still won't know who he is, and where he is, no matter how much knowledge he piles up, until that moment when he is ready to say,

This place is aisle five, floor six, of Main Library, Stanford.

I am Rudolf Lingens.³¹⁰

Here, Lingens learns all that there is to know about Lingens. He might read about Lingens' current location, interests, physical features and so on. None of this information will help him until he realises that he is Lingens. Without this additional belief, that he might express by saying “I am Lingens”, Lingens cannot know that the person he is reading about is himself. (Of course, if he learns enough, he can make the inference that he is Lingens, but it seems that any non-*de se* attitude will not help him. The thought “Lingens is Lingens” would not help him).

³⁰⁹ Recanati (2017: 180) defends a way of inferring, from the premises that **a** is *F* and the premise that **b** is *G* to the conclusion that something is both *G* and *F* without invoking a premise to the effect that **a=b**. This is referred to as trading upon identity. However, this move requires that two mental “files” be associated with one another, in that they both fall under a third file based on multiple epistemically rewarding (ER) relations. It seems that the self would have a file that is based on multiple ER relations. Even if the self file (presumably the relevant one for *de se attitudes*) was not based on multiple ER relations (which seems implausible), it seems that trading on identity still requires some association between two files, and this may be enough to maintain Perry's point that there needs to be some sort of belief to the effect that “I am Perry”.

³¹⁰ Perry 1977: 492

Perry also discusses some responses that one might make to his argument for *de se essentialism*. One is that the “I” is only a communicative shortcut.³¹¹ So “I” may simply be shorthand for some other proposition which picks out the bearer uniquely. One example could be “The person having this (token) thought”. However, it seems that this strategy will only work on the assumption that I am the person who is having this thought.³¹² Otherwise, having that thought is unlikely to affect their behaviour. (They may also have an unusual psychology, but that seems to be irrelevant). This point is also supported by the Lingens case.

Some further supporting arguments are made in favour of *essentialism*. Attempts to replace “I” with a description fail for two reasons. One is that these descriptions explain the ‘...action only on the assumption that I believed that I was the only [one who met the description], which brings in the indexical again’.³¹³ Perry may try to replace an indexical with a description of himself, such as the only philosopher in the supermarket. However, for this strategy to work Perry must believe himself to be this philosopher in the supermarket. The second reason is that the description may not be accurate. If so, ‘the proposition supplied by this strategy would be false, while what I came to believe, *that I was making a mess, was true*’.³¹⁴ For instance, Perry might not be the only philosopher in the supermarket.

These arguments suggest that there is something essential about the use of “I” in describing the subject’s beliefs. If we remove the “I” then we lose the ability to explain the agent’s behaviour. The point also seems to apply to the agent’s own beliefs as well. If they did not have some equivalent of the natural language “I” in thought, then it seems that they also lose the ability to represent certain properties as belonging to them. This suggests that *de se essentialism* applies not only to our

³¹¹ Ibid 368

³¹² Ibid 368

³¹³ Ibid 369

³¹⁴ Ibid 369

descriptions of agent's beliefs, but to their beliefs as well. So, it seems that having a *de se thought* is essential to grasping propositions in the correct way.³¹⁵

Lewis argues that some forms of knowledge cannot be had without a *de se thought*. Famously, Lewis considers the case of the two Gods who know all non-*de se* propositions but are still ignorant. This ignorance can only be removed by their coming to believe a *de se* content.

Consider the case of the two gods. They inhabit a certain possible world, and they know exactly which world it is. Therefore they know every proposition that is true at their world. Insofar as knowledge is a propositional attitude, they are omniscient. Still I can imagine them to suffer ignorance: neither one knows which of the two he is. They are not exactly alike. One lives on top of the tallest mountain and throws down manna; the other lives on top of the coldest mountain and throws down thunderbolts. Neither one knows whether he lives on the tallest mountain or on the coldest mountain; nor whether he throws manna or thunderbolts.³¹⁶

Lewis argues that for the Gods to know which of the two they are, they will need to come to believe a *de se* thought. For Lewis, this means that they will need to make use of what he calls a 'centred possible world', rather than a proposition, though this technical notion will not be important in what follows I will give a brief description here. A centred possible world is a world paired with a space-time point that corresponds to the individual.³¹⁷ So, to entertain certain sorts of knowledge, we need to make use of *de se contents*, on Lewis's view.

I will make use of Perry's model of *de se attitudes* according to which 'Anyone at any time can have access to any proposition. But not in any way'.³¹⁸ Having an appropriate belief state is necessary for having a *de se attitude* as without that one will not realise that the content is about them. I will argue

³¹⁵ Ibid 375

³¹⁶ Lewis 1979: 520-521

³¹⁷ Ibid 531

³¹⁸ Perry 1979/2010: 375

that these states can be well understood as requiring an indexical component in the thought that is tokened when entertaining these beliefs in section 3.3. So, one has a *de se attitude* when one tokens an indexical thought about themselves and that indexical functions similarly to the English “I”.

Together, Perry and Lewis have made a powerful case for *de se essentialism* on the basis that we cannot explain certain agents’ actions without making use of *de se attitudes* and that we cannot entertain certain kinds of knowledge without making use of *de se attitudes*. These arguments suggest that having a *de se thought* is essential to grasping propositions in a self-aware way, which is claim (i). In section 3.2.3 I will consider some more contemporary arguments in favour of *de se essentialism*, particularly Ninan’s (2015) defence of *de se essentialism*. This is to respond to *de se sceptics*, in particular Cappelen and Dever (2013).

3.2.1 Expanded Doctrine of Propositions

In this section I will give a more detailed defence of claim (i), that having a *de se thought* is essential to self-awareness. Some of the key motivations for this view have been considered above, but there is more to the argument than the examples given by Lewis and Perry. These examples are meant to show that *de se attitudes* create a unique problem for theories of propositions which marks them out as essential or indispensable. *De se essentialism* is taken as an issue for theories of propositions in general, rather than for a particular theory of propositions. Here I will give the account of propositions they cause a problem for.

This is because the problem of *de se attitudes* can arise for the *expanded doctrine* of propositions. This doctrine is a set of desiderata for propositions. Importantly, ‘[f]or all they say, contents might be Fregean propositions, Russellian propositions, or sets of possible worlds. If we can find a problem for the expanded doctrine, then we have found a problem for anyone... regardless of their other views about the nature of propositions.’³¹⁹ The expanded doctrine is composed out of five features that are intended to be neutral between particular theories of propositions. These are the following features.

³¹⁹ Ninan 2016: 102

(1) Two-Place Relation: An attitude relation is a two-place relation between an agent and a content.

(2) Frege's Constraint: Contents are assigned to attitudes in a manner that accommodates Frege cases.

(3) Absoluteness: the contents of attitudes do not vary in truth value across individuals or times.³²⁰

(4) Agreement: Agreement is a two-place relation between a group of individuals and a content (more detail below).³²¹

(5) Explanation: Where an agent X's action A are explained by his having set 1 of beliefs and set 2 of desires then an agent Y with set 1 and set 2 will, other things being equal, also perform A.³²²

Most of the arguments in favour of *de se essentialism* have focused on the first three of these features. For instance, Perry mentions them explicitly.³²³ However, these approaches to supporting *de se essentialism* have been challenged by *de se sceptics*. In the following section I will present some of the arguments made by *de se sceptics* before considering responses to them.

3.2.2 The Challenge to *De Se Essentialism*

De se essentialism has been challenged by Cappelen and Dever (2013) who can be described as *de se sceptics*. They claim that there is no unique problem posed by *de se attitudes* and therefore no special *de se thought*. Instead, they claim that 'when we look at the interaction between indexicals and opaque contexts, there's just the general problem of Frege puzzles.'³²⁴ If Cappelen and Dever are correct then there is no tension within the expanded doctrine, it just reinforces the second feature, that propositions be able to accommodate Frege cases. If *de se thoughts* really are no different from other kinds of thoughts, then we can exchange them with other thoughts as easily as we could with

³²⁰ Ibid 92 The first three are from Perry 1979/2010: 367-368

³²¹ Ninan 2016: 100

³²² Ibid 102

³²³ Perry 1979/2010: 367-368

³²⁴ Cappelen and Dever 2013: 59. Other cases are taken to be irrelevant (Heimson/Hume) or to be question begging or otherwise problematic (as in Lewis's case of the two gods). See Stalnaker 1981: 144, Cappelen and Dever 2013: 89. See Cappelen and Dever 2013, chapter five, for extended discussion.

any other pair of co-referring thought. The evidence of there being something special about *de se thoughts* stems from the cases that *de se essentialists* have presented.³²⁵ If these cases are undermined then *de se thoughts* require no special treatment. Cappelen and Dever raise several distinct lines of criticism.³²⁶ Here, I will focus on their reaction to Perry's cases as his is the model of *de se attitudes* that I am adopting here.

Cappelen and Dever replicate some alleged examples of *de se essentialism* and create analogous versions without making use of *de se thoughts*. In which case, it is not special a phenomenon but a specific kind of Frege case. One is the case of Lingens the amnesiac. Cappelen and Dever claim that by changing the wording of the case from Lingens knows *who* to knows *that*, they can give the following Frege counterpart:

We are trying to locate Superman. We are in the Stanford library, and read a number of things in the library, including a biography of Clark Kent. We believe any Fregean thought you think might help us. We still won't know that Clark Kent is Superman no matter how much knowledge we pile up, until that moment when we are ready to say, "Superman is Clark Kent."³²⁷

In so far as this superman case is analogous to the case of Lingens it appears both are just Frege cases. So, the Lingens example no longer seems to show us that there is anything special about *de se attitudes*. Instead, it is just a case of having an informative identity statement. Namely, that "I" is "Lingens". Similar Frege counterparts are given for the other cases discussed in 3.2. This presents a challenge to *de se essentialists* to have an argument to the effect that there is a unique problem for theories of propositions that stems from *de se attitudes*.

³²⁵ Perry 1977, Perry 1979, Lewis 1979.

³²⁶ 'Perspectival representations are representationally essential: perspectival makes available ways of representing the world that are not available with nonperspectival devices' (Cappelen and Dever 2013: 1).

³²⁷ Cappelen and Dever 2013: 63

3.2.3 The Problem Posed by *De Se Attitudes*

Ninan takes up this challenge, by attempting to show that *de se attitudes* cause a unique tension between Agreement and Explanation (defined below). If Ninan is correct, then the problem posed by *de se attitudes* is not simply a Frege case. Instead, it makes the expanded doctrine false and shows why *de se attitudes* present a problem of their own. I will begin by elaborating on Agreement and Explanation (I use capitals to mark these as technical terms).

‘Agreement is a two-place relation between a group of individuals and a content. A group of individuals stands in the agreement relation to a content p iff all the members of the group believe p . Let B_p^x and B_q^y be token beliefs belonging to x and y respectively. Then x and y agree on something in virtue of x ’s having token belief B_p^x and y ’s having token belief B_q^y iff $p = q$. AGREEMENT only concerns beliefs.’³²⁸ Less robotically, x and y Agree when they have beliefs that are equivalent to one another. For instance, if each is in a belief state such that they believe that the sky is blue, then they Agree. If one believes that the sky is not blue, then they do not Agree. However, there is an understanding of Agreement which concerns their desires as well. These will be cases in which two people have the same desires, relative to the case at hand.³²⁹ So if we are watching a sports match and desire that the same team win, we have Agreement in terms of our desires relative to that case.

The second part of the expanded doctrine is Explanation which goes as follows, ‘Suppose the fact that x performed action α is explained by the fact that x has beliefs $B_{p_1}^x, \dots, B_{p_n}^x$ desires $D_{q_1}^x, \dots, D_{q_k}^x$. Then, if y has beliefs $B_{p_1}^y$ and $B_{p_2}^y$ desires $D_{q_1}^y$, then, other things being equal, y will also perform α .’³³⁰ The conditions for Explanation are met when two agents Agree, so their beliefs (at least their relevant ones) are equivalent, and when they want the same outcome for a given event. This means that, all other things being equal, they will perform the same action as one another, as they presumably both believe that it is the optimal means to their end, and they want the same end.

³²⁸ Ninan 2016: 100

³²⁹ Ibid 100

³³⁰ Ibid 102

Ninan shows that (4) and (5) can be inconsistent and this is the problem that we get from *de se attitudes*. Ninan outlines his problem as follows:

‘AGREEMENT would seem to entail if two agents agree on how things are and on how they would like them to be, then those agents have the same beliefs and desires, i.e. they believe and desire the same propositions. And EXPLANATION tells us that if two agents have the same beliefs and desires, then, other things being equal, they will behave in the same way. But then it follows from those two claims that if two agents agree on how things are and on how they would like them to be, then, other things being equal, they will behave in the same way. But, as I shall now argue, this final claim is false. If my argument is sound, it follows that AGREEMENT and EXPLANATION are inconsistent, and that the expanded doctrine is false’.³³¹

Ninan’s aim seems to be to show that *de se attitudes* can create cases in which, if Agreement is true, then Explanation is false. That is, if two agents Agree, in that they have the same beliefs and desires, relative to the case, then if Explanation were true, they would perform the same action. But they do not perform the same action. So, it seems that Explanation isn’t true when Agreement is. In which case, we can either accept Agreement, so that Explanation is false or claim that Agreement is false, but Explanation is still true. Either way, a part of the expanded doctrine is false.

Ninan’s argument for this claim uses a thought experiment in which Ninan is being chased by a bear, and you are watching this chase. Both you and Ninan have the belief that Ninan is being chased by a bear and the desire that Ninan not be killed by the bear. Both you and Ninan have the belief that the best way to avoid being killed by a bear is to curl up in a ball. So, to avoid death by bear, Ninan curls into a ball.

³³¹ Ibid 103

‘But now we have a problem. For we have:

(a) The fact that I curled up into a ball is explained by the fact that I have beliefs B^m_{p1} and B^m_{p2} and desire D^m_{q1} .

And we have:

(b) You have beliefs B^y_{p1} and B^y_{p2} and desire D^y_{q1} .

But from (a), (b), and EXPLANATION it follows that, so long as other things are equal, you too will curl up into a ball. But this prediction seems false’.³³²

This is a problem caused by *de se attitudes*; it seems that one of the distinguishing features of Ninan’s belief is that *he* believes that *he* is being chased. So, he believes I AM BEING CHASED BY A BEAR, and that curling up into a ball is the best way to avoid being mauled by a bear. Watching, you may think NINAN IS BEING CHASED BY A BEAR and believe that curling into a ball is the best way to avoid being mauled by a bear. You both desire that Ninan not be mauled by a bear. So, you and Ninan have met the conditions for Agreement. It follows that you should both perform the same action because of Explanation. So, you will both curl into a ball. But this prediction seems false. Ninan will curl into a ball; you will likely do something else. It follows that either Agreement or Explanation is false. From which it follows that the expanded doctrine is false. This is the unique problem of *de se attitudes*.

3.2.4 Defending the Problem of the *De Se Attitudes*

In this section I will discuss three objections that might be made to Ninan’s argument for *de se essentialism*. The first is a consideration of whether all things are equal. If not, then the conditions for Explanation are not met and there is no conflict. The second is that Ninan’s problem is still just a Frege case and so causes no unique problem. The third is that there is no Agreement in Ninan’s case, so there is no conflict between Agreement and Explanation.

³³² Ibid 104. Here the change from B^x_{p1} to B^m_{p1} seems to indicate that this is a belief of Ninan’s, which he might express as “my belief that p1”.

3.2.4.1 *Are All Things Equal?*

First: Are all things equal here? This is important as it is a requirement of both Agreement and Explanation that all else be equal. Perhaps for all else to be equal here you must also be being chased by the bear. In which case the thought experiment Ninan describes is not one in which all else is equal, so it is not one in which neither Agreement nor Explanation can be said to hold. So, they cannot contradict one another either. However, this misses the point. Both you and Ninan will have belief tokens that are equivalent for this case. That is, both you and Ninan agree that Ninan is being chased by a bear. Otherwise, you do not agree with Ninan that he is being chased by a bear.

Perhaps things are unequal as the actions that you and Ninan can perform differ. Cappelen and Dever deny that all is equal by claiming that different actions are available to be performed by different agents. Whilst two agents can be motivated to perform the same action, they cannot in fact perform the same action.³³³ So, they claim that: ‘Someone other than François, e.g. Dilip, can be rationally motivated to perform the action that François climb into the tree. Dilip just can’t perform the action’.³³⁴ The reason that they cannot both perform the same action is that the action is specific to its being François that climbs the tree. The view here is one on which actions are typed according to the agent performing them as well as by the kind of action. The action of Dilip climbing into the tree is not one that François can perform, even if François is near a tree and is a capable climber. So, for all things to be equal, there needs to be the same agent considered in each case. If so, then all things are not equal and the conditions for Explanation are not met, so it cannot conflict with Agreement.

However, it seems that agents can perform the same action and we can easily imagine that both agents are able to curl into a ball or climb a tree.³³⁵ The response by Ninan seems to amount to a rejection of Cappelen and Dever’s method of counting actions by correlating them with an agent who

³³³ Cappelen and Dever 2013: 52

³³⁴ Ibid 52. Here climbing the tree plays a similar role to curling into a ball.

³³⁵ Ninan 2016: 106

performs that action. If they do maintain this method of counting actions then their view rules out the possibility that two agents can perform the same action, which seems absurd.³³⁶ At least, it seems to require an argument to the effect that we cannot type actions in an agent neutral way. In the absence of such an argument it seems that there is a sense in which can maintain that all else is equal. Whether this understanding of action is something that Cappelen and Dever maintain throughout is uncertain, but if they do not then it seems that they should weaken their case against *de se attitudes*, as there no longer seems to be a reason to deny that the cases are equal.

3.2.4.2 *Isn't this a Frege Case?*

Second: Is this problem unique to the *de se attitudes* or is it possible to have a similar Frege case?

Ninan argues that such a case is implausible because in cases in which there is Agreement about non-*de se beliefs*, it seems plausible that these beliefs will meet Explanation as well.³³⁷ To illustrate, suppose that we are in the woods and see someone being chased by a bear. I believe that CLARK KENT IS BEING CHASED BY A BEAR, and that if he curls into a ball he will not be mauled. I also desire that he not be mauled. Suppose that when you see this person being chased by a bear, you believe and desire the same things, except that you believe SUPERMAN IS BEING CHASED BY A BEAR. So, we might try to tell this person that they should curl into a ball. In which case, there is no conflict between Agreement and Explanation. In which case, we do not have an instance of Ninan's problem. A response is that this is not the correct prediction. If you believe that Superman is being chased by a bear, whilst I am not aware that Clark Kent is Superman, then we might offer different advice. You may suggest that he simply flies away, whilst I suggest he curls into a ball. Isn't this an instance of the same problem recast as a Frege case? To be a Frege case, it must be the case that if Agreement is true then Explanation is false.³³⁸ Here, Ninan might deny that we have Agreement, as the two sets of

³³⁶ Ibid 106

³³⁷ Ibid 109 'For if the elements of a non-*de se* pair are agreement similar, there is a case to be made that they will also be functionally similar'.

³³⁸ As Ninan's case is intended to show an inconsistency between Agreement and Explanation, so it can be expressed as the following biconditional: $\sim (A \leftrightarrow E)$. However, the focus should be on $A \rightarrow \sim E$, because if the conditions for Explanation are met then it follows that the two agents Agree.

beliefs are not identical. I do not believe that the person being chased is capable of flight, whilst you might. But then the onus is on Ninan to say why this case is not Agreement whilst in *his* case there is Agreement. Presumably, the difference will be in knowing what the agent can do. So, they do not have the same set of beliefs, for instance, that that man can fly. If so then they do not Agree, so the conditions for Explanation are not met, so there is no inconsistency. So, this is not an instance of Ninan's problem.

However, I may also believe that superman can fly. It seems that we then Agree on this, so we should agree that that man can fly. I am prevented from Agreeing with this statement because I do not believe that Superman is Clark Kent. So, I do not believe that Clark Kent can fly. Whilst this additional belief makes the case a Frege case, it is not an instance of Ninan's problem. This is because we do not meet the conditions for Agreement and the same holds for Explanation. So there does not seem to be any inconsistency here between Explanation and Agreement.

Whilst this discussion does not prove that there cannot be a Frege interpretation of Ninan's case, it does show that it would be difficult to do so. This is because when Ninan's case becomes a "standard" Frege case we lose the inconsistency between Agreement and Explanation. Either we do not Agree, because we have different beliefs about the abilities of the agents, or we Agree and perform the same action. Either way, there does not seem to be an inconsistency between Agreement and Explanation when we have only a Frege case.

3.2.4.3 No Agreement in Ninan's case

A third objection is that there is not Agreement in Ninan's case. Does the observer have the same belief as Ninan does, if one is expressed indexically (Ninan's) and one is not (the observer's)? Ninan assumes that there is Agreement. 'Now if, as I am assuming, we count as agreeing that I am being chased by a bear in virtue of my having this *de se* belief and your having the corresponding *de te* belief, then it would seem to follow from AGREEMENT that these beliefs have the same content.'³³⁹

³³⁹ Ibid 103

This assumption is perhaps something that Ninan's opponents would want to deny. If Ninan and the observer do not Agree, then Ninan's counter argument does not work. There is no Agreement and therefore no Explanation either. So, there is no contradiction between the two.

However, if one adopts the perspective of the sceptic, then it appears there should be Agreement, since for the sceptic there is no substantive difference between Ninan's belief and the observer's belief. There is nothing special about the *de se attitudes* on the sceptic's view. The thought that "I am being chased by a bear" could just as easily be replaced by the thought "Ninan is being chased by bear" without a change in explanation. In this case, the two agents in Ninan's case Agree with one another. But then they should act in the same way, in virtue of Explanation. This seems implausible. So, arguing that Ninan and his observer do not Agree seems to undermine the sceptical position.

If *de se essentialism* is correct, then Ninan's belief state and the observer's belief state can have the same content. Where they differ is in terms of how they are accessed. This also has the advantage of preserving a level of content at which *de se attitudes* and non-*de se attitudes* express the same content.

In this section I have aimed to provide a convincing argument for the truth of (i). For many, the truth of (i) means that there should be some indexical element to these attitudes. This is a claim that an indexical element, such as "I" is essential to expressing certain attitudes. I will introduce and defend this claim in the following section.

3.3 Claim (ii)

Claim (ii) is that these *de se thoughts* make use of a context-sensitive character to determine their reference. This claim is often taken to be a part of *de se essentialism*, and often it seems to be assumed that a *de se thought* will involve an indexical, as this seems to be essential to their being represented in natural language. However, it is not clear whether this should also be a requirement at the level of thought and there are positions that accept (i) without accepting (ii). Notable examples include Christopher Peacocke, Gareth Evans and John McDowell who argue that reference is

determined by a (*de re*) sense instead of by character.³⁴⁰ In this chapter I will defend (ii) but I will not engage too much with these views as my aim is to show how the CTM can be compatible with indexicals, rather than to show that (i) entails (ii).³⁴¹ (ii) is significant as it allows for the possibility of context-sensitive thoughts on a CTM.

The indexicality that I have in mind is Kaplan's context-sensitive characters, whose contents are changeable, I will call these "changeable" characters to distinguish them from cases which are not context-sensitive in this way. He claims that 'Indexicals have a *context-sensitive* character. It is characteristic of an indexical that its content varies with context'.³⁴² The important point here is that these indexicals do not function in a comparable way to other singular terms. They do not have a fixed reference across all possible contexts of use, rather their reference depends on the context of use in which the indexical is tokened. It should be noted here that though *de se attitudes* might be indexical in Kaplan's sense, they may not actually change their referent in many, if any, contexts of use. The *de se attitudes* will not be practically indexical.³⁴³ Practical indexicals are ones that change their reference across contexts of use. This does not mean that they are not indexical at all. They may still depend on the context of use to get a reference. I will argue that something like (ii) is useful for accounting for claim (i). That *de se attitudes* are essential chimes well with them requiring indexical thoughts, as I will argue below.

My motive for selecting Kaplanian character as a way of understanding indexicality is partly to keep the debate in line with opponents of *de se essentialists*, such as Cappelen and Dever who '...take Kaplan's account of indexicality to be the most philosophically minimal and conservative account since it preserves a level of content at which there is no difference between indexical and non-

³⁴⁰ See Peacocke (1981: 189 – 90), Evans (1982: 35, 159, 205), Evans (1981: 280, 284) and McDowell (1984b: 288 – 9).

³⁴¹ If (ii) were false then it would be less likely that there are thoughts that exhibit character underdeterminacy, but it may not be fatal to this position.

³⁴² Kaplan 1989: 506

³⁴³ Ball 2015: 354

indexical expressions.³⁴⁴ So the motivation for choosing Kaplanian indexicalism to model *de se* thoughts is to avoid making too many commitments for the *de se sceptics* to object to.³⁴⁵

Together (i) and (ii) give us the claim that to grasp certain representations about ourselves we need to have a context-sensitive, *de se thought*. So, if we are to grasp certain propositions about ourselves then we need to make use of an essential indexical. In which case, we would have a context-sensitive thought. Claim (i) is made very compelling by Perry's examples. This provides an argument for Character-Underdeterminacy of thoughts, where Character-Underdeterminacy can be understood as follows:

Character-Underdeterminacy: A structured representational item S is character-underdetermined if and only if a component of S makes an explicit reference to the context to establish content.

The *de se thoughts* will contain an indexical element which makes explicit reference to the context of use to establish its content.

Claim (i) and (ii) together also seem to mean that these indexical thoughts are foundational, in the sense that they are of '... a level of thought which is fundamental in this strict sense: every thought about Gs which is not of this level is conceived to be made true by the truth of thoughts which are of this level'.³⁴⁶ If they are replaced by a different thought then there has been a loss, as per *de se essentialism*. So, they cannot be made true by any non-*de se* thoughts. This makes them the kind of foundational thought that we are interested in investigating (see chapter 1).

³⁴⁴ Cappelen and Dever 2013: 10

³⁴⁵ There are other accounts of "I" that are available that I could make use of here. For instance, De Gaynesford (2006) argues that the word "I" is a deictic term and not a pure Kaplanian indexical. However, this account is orthogonal to this discussion of thoughts. Furthermore, if the mental "I" should also turn out to be deictic that would not hurt the claim that thoughts can be character-underdetermined, though it would require making sense of how one remains the salient referent of this term. On other accounts there is an epistemically rewarding relation between the concept and the object (Recanati 2012) or as a concept we acquire in virtue of phenomenal experiences (Guillot 2016). An exhaustive discussion of these alternatives would go beyond the scope of this chapter, but the alternatives serve to put pressure on the indexical view considering Millikan's objections.

³⁴⁶ Evans 1982: 112

3.3.1 Defence of claim (ii)

(ii) is the claim that *de se thoughts* are indexical, in the sense that they have a context-sensitive (Kaplanian) character, where '[i]t is characteristic of an indexical that its content varies with context'.³⁴⁷ The following is a summary of the indexicalist position:

The claim that indexical modes of presentation cannot be captured by means of objective, non-indexical descriptions is the claim that, to the extent that they can be captured by descriptions, the descriptions in question will *themselves* contain indexicals. The indexical component in them is therefore ineliminable or "essential": any attempt to cash it out descriptively will produce an indexical residue...³⁴⁸

This is the view that there is a class of representations that cannot be replaced by a non-indexical component without there being some loss. *De se indexicalism* is a more particular version of this claim, which states that thoughts about oneself cannot be replaced by a non-indexical symbol without a change in what that thought can explain. In this section I will defend the claim that *de se indexicalism* and that this indexicality contributes to *de se thoughts* being essential.

Often, it is assumed that if (i) is true then so is (ii). This is not an obviously true claim however and needs defending. Here I will argue that something like (ii) seems necessary to keep something special to *de se attitudes*, or to explain why they pose a unique problem for accounts of propositions. This is to avoid the *de se thoughts* collapsing into a Millikan type view.³⁴⁹ On the Millikan type view, there are no indexical elements in thought. However, if such a claim is true then it seems that we will struggle to explain why (i) is true. If the concept that we have for the self is a proper name, then it is not obvious why this should pose a unique problem for accounts of propositions. Instead it would be a Frege case in which one needs to realise that e.g. Millikan is @RM. This would not be a unique

³⁴⁷ Kaplan 1989: 506

³⁴⁸ Recanati 2012: 33

³⁴⁹ A similar objection is attributed to Papineau (2006): 'Papineau therefore objects to the view that that such concepts are "demonstrative", and work like indexicals. They are, he maintains, more like proper names' (Recanati 2012: 76).

problem for propositions to answer to. In section 3.3.2 I will elaborate on Millikan's view but then argue against it in 3.3.3. I will consider an indexing view that is also incompatible with (ii).

3.3.2 Millikan's View, the Case Against (ii)

Millikan is critical of the idea of indexical elements in thought, as is illustrated in the following quote:

Do I succeed in identifying the content of various tokens of my mental "@RM," that is, do I succeed in reidentifying myself, only because I grasp for each token of "@RM" independently that it bears a certain adapting relation me? Isn't it more reasonable that my mental "@RM" is simply a mental proper name? I take different tokens of "@RM" to refer to the same [individual] not because of their individual contexts, not as a result of some relation each of these tokens independently bears to me, but simply because they are tokens of the same type.³⁵⁰

Here, Millikan expresses her disapproval of the idea of an indexical element in thought. She finds it implausible that we would have a context-sensitive token when a name would seem to fulfil the same task of referring to oneself without the need to take the context of use into account. In which case, we should reject (ii). On the view that she presents (though does not ultimately endorse) '...my mental "I," my "@RM," is not an indexical' but a mental proper name.³⁵¹ She describes this as a Millian name so that the only semantic contribution it makes is its referent.³⁵² So, she opts for the view that either the self does not need to be represented at all or if it does then it may only need to be a (Millian) name.³⁵³

³⁵⁰ Millikan 1990: 11

³⁵¹ Ibid 11

³⁵² Cumming 2016: S2.1

³⁵³ Millikan 1990: 12, she writes '...my intentions are not designed to guide anyone's actions but my own. Hence, they have no need to explicitly represent me. I do not have to take into account variations in whose head a token of "@RM" appears, nor variations in whose action it is supposed to guide. But, once again, this inarticulateness in how the self is represented has nothing to do with indexicality'. This leaves it open that the self is represented inexplicitly, and there may be a parallel with Perry here, whose view I will discuss in chapter 4. At present, it seems as though Millikan opts for any inexplicit representation to occur by default. There is no one else to whom this representation could apply.

The challenge that Millikan presents here can be put as follows. Why should we prefer (non-practical) indexicality when we could make use of a proper name for ourselves? Here (as noted above) a practical indexical is one that ‘...will very often link different tokens to different referents, even if we hold the speaker fixed. Call terms and mental files that shift reference in this way practically indexical’.³⁵⁴ Non-practical indexicals will not shift their referents. The bite of Millikan’s objection comes into focus when we consider that having non-practical indexicality seems to come with additional commitments that we are not certain we should be making. For instance, that there might be some circumstances in which it could become practical indexicality e.g. if my “self” symbol were to appear in your head. It also seems to add an additional mechanism of saturation, where saturation is the process that ‘... takes place whenever the meaning of the sentence includes something like a “slot” requiring completion or a “free variable” requiring contextual instantiation’.³⁵⁵ That would involve assigning a value to the indexical. So, then the question is what warrants positing this additional complication.

3.3.3 Responses to Millikan

One response is to say that these *de se attitudes* are only properly expressed in natural language by using indexicals, such as “I”.³⁵⁶ Even if this true, however, it does not seem to follow that the thoughts that we have must also use such an indexical component. That seems to suppose that natural language and thoughts are analogous in, at least, this respect. That begs the question.³⁵⁷ We must assume that thoughts are like natural language, but this is what we are trying to find out. This argument will not do.

³⁵⁴ Ball 2015: 354

³⁵⁵ Recanati 2004: 7. If one takes a Kaplanian view then saturation we must still find a content for that character. However, if one wished to maintain a Kaplanian account of saturation, on which the character interacts with a context of use to provide a content without changing the representation, then the indexical element would not need to be replaced.

³⁵⁶ This is denied by some defenders of the *de se* e.g. Bermudez (2017: 12) who claims that we can use our own names as deferred uses of “I”.

³⁵⁷ Ball 2010 makes a similar point, to which I am sympathetic.

A point against Millikan's understanding of *de se attitudes* is that in the case of a name, it seems that it must be possible for the name's reference to be forgotten, or the name may fail to refer, unlike in the case of *de se attitudes*. Millikan's name would be somewhat different, in that it seems very difficult to imagine cases in which one has a *de se thought* that fails to refer.³⁵⁸ In fact, tokened *de se thoughts* seem to be incapable of failing to refer. If someone can token them, then they can refer to themselves with it, and have done so.³⁵⁹ The same does not seem to be true of names. This seems to be so regardless of the theory of names that one picks. These disanalogies between the way *de se thoughts* work and the way that names work gives us reason to resist Millikan's argument.

For this point to be convincing it is necessary to prove that "I" can have these features as well. It seems that it can. Tokening "I" seems to make it refer.³⁶⁰ It is hard to conceive of a tokened mental "I" failing to refer. If it is tokened (in a use rather than mention context), then it will refer to the thinker that tokened it. For these reasons it seems that we can justify making use of an indexical account, rather than one of names, to account for *de se attitudes*.

A counter example might be cases where I quote, or think of, someone saying "I am brilliant" who does not really exist.³⁶¹ One way to adopt this view is to be an anti-realist about fictional entities, such as Sherlock Holmes. Proponents of such a view include Corozza and Whitsey (2003), who claim that in such cases there is a deferred reference from the actor/author to the fictional entity.³⁶² However, there is no fictional entity (on the grounds of parsimony), so the utterances that are attributed to these characters do not express propositions (or if they do they are merely token-

³⁵⁸ Anscombe 1975: 22-23

³⁵⁹ Ibid: 28

³⁶⁰ There are several ways of understanding this referring. On one understanding the reference cannot fail at all, on another the act of tokening the symbol is sufficient (de Gaynesford 2006: 4, 165). Which is the best way of understanding this notion will not be too significant in the following discussion.

³⁶¹ There is a related worry here in how to understand the role of "I" in natural language, for instance, whether it refer to the producer or the user of the term (De Gaynesford 2006: 43-44). This is significant as "I" is often given a character along the lines of referring to the producer or the user of the term, but these cases are not equivalent.

³⁶² Corozza and Whitsey 2003: 135

reflexive propositions).³⁶³ So if one takes such a view it is not necessary that “I” refers to anyone or anything.

However, we are certainly not forced into an anti-realist position. Realist positions, according to which fictional entities are abstract objects, are also available. A sketch of such a view goes as follows: ‘...wholly fictional characters like Sherlock Holmes, though real, are not real people. Neither physical objects nor mental objects, instead they are, in this sense, abstract entities. They are not eternal entities, like numbers; they are man-made [artefacts] created by fiction writers. But they exist just as robustly as the fictions themselves, the novels, stories, etc. in which they occur’.³⁶⁴ On this view, indexicals can refer even when used in fiction. If a view along these lines were correct, then there would not be a counter example to “I” referring whenever it is used.

This realist approach is not necessarily better, however. Now there is a commitment to the claim that artefacts can refer to themselves. This is strange as the artefacts themselves do not do anything.

Mostly they are just used to do things. So, whilst there is reference on a realist view, it should not be treated as self-reference. However, as there is no agent, it seems that artefacts from fiction do not prove to be a counter example to the claim that whenever an agent uses “I” it can refer. This is unlike the case of names. A similar response can be made to anti-realist responses as well.

Outside of cases of indexicals in fiction, it seems that “I” refers, and these fictional cases seem significantly different from cases in which one uses the token for oneself. Having the property of deferred reference seems to make the fictional cases particularly different from cases in which one refers to oneself. It is also hard to imagine a tokened mental “I” failing to refer. If it is tokened, then it will be referring to that thinker. For these reasons it seems that we can justify making use of an indexical account, rather than one of names, to account for *de se thoughts*.

³⁶³ Ibid 134 – 135

³⁶⁴ Salmon 1998: 293

This argument gives us some reason to accept that *de se attitudes* cannot be properly accounted for by names, contra Millikan. This goes some of the way to supporting the claim that non-indexical forms of reference would struggle to have the features that we need to account for *de se attitudes*, which supports (ii).

3.3.4 Indexing *De Se Attitudes*

Another alternative to (ii) is that rather than having a SELF concept whose reference is determined each time it is tokened, we instead index the concept to fix its reference.³⁶⁵ So we may have several demonstratums, which lets us list the demonstratums as ‘first demonstratum, second demonstratum, . . . (some of which may be null) as features of a context. We then attach subscripts to our demonstratives and regard the n-th demonstrative, when set in a context, as rigid designator of the n-th demonstratum of the context. Such a rule associates a character with each demonstrative’.³⁶⁶ To give an illustration, this is the difference between having a case such as “I” and “I_{NN}”, where the latter has been indexed to an individual, NN.³⁶⁷ This view avoids the worry that the names view faced because it seems that indexing requires either a perceptual or anaphoric link to the object that it refers to. They will be linked to their referent in *de se attitudes* via the agent’s awareness of themselves and it seems that they won’t fail to refer as names do. But these indexed components also do not seem to be indexical anymore.³⁶⁸ This indexing view relinquishes (ii).

There is, however, an issue with this approach. We do not have a guarantee that one is aware that the object that is indexed with “I_{nn}” is oneself. To see this, consider again the case in which Perry is searching for the messy shopper, except that there are mirrors at either end of the counter so that as he pushes his trolley down the aisle he sees a reflection of the messy shopper going down the aisle on the other side.³⁶⁹ Perry points and says “I believe that he is making a mess”.³⁷⁰ Suppose that here

³⁶⁵ I’d like to thank Dr Nat Hansen for this point, in discussion.

³⁶⁶ Kaplan 1989: 528

³⁶⁷ Peacocke (1981: 189 – 90) seems to endorse a similar view.

³⁶⁸ Ball 2010: 14

³⁶⁹ Perry 1979/2010: 371

³⁷⁰ Ibid: 371

we have a case in which the “he” is indexed to become “he₁”. Perry is not aware that he has indexed himself as the referent of “he₁”. This example shows that indexing oneself is not sufficient to realise that one is referring to oneself.³⁷¹

For the indexing account to be a useful account of *de se attitudes* one will presumably have to know that they are the NN. In which case, they seem to have to judge that “I am the NN”. Doing this involves a non-indexed first-person concept, the “I”. In which case, they must have already been able to think of themselves in a *de se* way prior to having the indexed first-person concept “I_{nn}”. So, the self-awareness of *de se attitudes* cannot be entirely accounted for by mental indices.³⁷² Attempts to do so only move the question. So, the prospects for an indexing view of *de se attitudes* do not seem to be promising.

3.3.5 Over Intellectualising?

A further response is that the indexical account makes the possession of indexical concepts too difficult as the function might require the possession of other quite complex concepts. In more detail the objection goes as follows:

The token-reflexive description [of the index model] identifies as the cognitive significance of the I-concept is a fairly complex conceptual representation. One might object that this commits the approach to an implausible form of intellectualism. A grasp of the I-concept, on this view, presupposes the prior grasp of further, fairly abstract concepts, at least the concept of a thinker and the concept of a thought; or, if one wants to paraphrase the token-reflexive description differently, concepts of an equivalent level of generality. But it is unclear whether a small child who has just come to master the I-concept, for instance, is also able to think of thoughts, or thinkers, as such.³⁷³

³⁷¹ Similar problems are possible in the case of names.

³⁷² I would like to thank an anonymous reviewer for their helpful comments in this regard.

³⁷³ Guillot 2016: 142

The challenge here is that the context-sensitive character of indexicals is often taken to be something like “the thinker of the current thought”. The worry is that we can acquire a concept that enables us to entertain a *de se attitude* prior to our acquiring these more complex concepts of THINKER or of THOUGHT. So, the concept that enables us to entertain *de se attitudes* must be simpler than the indexical concept. Otherwise we would not be able to have *de se attitudes* so early in life.

In response to this objection, it is not obvious that children have not acquired these concepts at an early stage in their lives. Recall Fodor’s understanding of nativism according to which encountering something is sufficient to cause us to have a concept of it (see section 2.4). If Fodor’s view is correct, then it seems plausible that young children do have these concepts. However, this does not entail that they have a great understanding of these concepts. They may have the concept without a great deal of encyclopaedic information to accompany it. Failing that response, it is also not obvious that children need to understand the mechanism by which their concept refers to have that concept or that a descriptive understanding of how an indexical “I” refers is correct.

If the above is correct, then this gives us some reason to accept (ii). It shows that there is an explanatory burden on accounts of the *de se essentialist* that cannot obviously be met by accounts that do not allow for some sort of (Kaplanian) changeable indexical element. In the next section I will discuss claim (iii).

3.4 Claim (iii)

(iii) holds that computations will interact with all instances of a given syntactic type in the same way.³⁷⁴ So whilst computations operate over particular instances of a syntactic type - they operate on tokens - they are not capable of treating tokens of the same type differently depending on the context in which they occur. This is because syntax is a local property of representations and computations can only track syntactic properties. Computations cannot consider non-local properties like the context of use. Computations can only pay attention to a semantic type indirectly, by paying

³⁷⁴ Ball 2010

attention to syntactic type. This is expressed as the idea ‘...that computation manipulates symbols in virtue of their formal syntactic properties rather than their semantic properties.’³⁷⁵ I will refer to this idea as the computation acting on a symbol *qua* type, rather than *qua* token. Treating a symbol *qua* token would involve paying attention to its non-local properties so that its semantic content is not just a product of its syntactic type. I will take a CTM to be committed to the claim that at least some of the processes in the mind are computations in this sense.

It is argued that exclusive sensitivity to local properties makes computation incompatible with context-sensitive indexicals, as context-sensitive indexicals need to be treated as tokens to get to the content that they express. It is possible that two different utterances of “I” will have different expressions, so in natural language we are treating them as tokens as well as types. It is this “treating as token” that will cause problems for the compatibility of these claims, as Argument 1 states (to be introduced below). This means that any indexicals in thought will be treated in the same way. In which case, all instances of a given indexical type will be treated as though they refer to the same object in all contexts. Argument 1 claims that this is implausible, so we should not think of thoughts as including indexical elements. (I will present Argument 1 below).

It seems to follow that if (iii) is true, then indexicals with a changeable character are inconsistent with the CTM. So, if the CTM is true then there cannot be mental indexicals. It follows that if we do have mental indexicals then the CTM is false, by contraposition. It is this inference, from (iii) to the claim that if we have mental indexicals then we cannot have the CTM that I will attempt to show is flawed in section 3.5. I will argue that this argument requires an overly strong conception of the CTM, opting instead for a view on which not all the mind needs to operate computationally. In this way it is possible to have both the CTM and mental indexicals.

One might think that the position advocated by Argument 1 here is implausible as we have indexicals in natural languages. In which case, there needs to be a way for computations to process these

³⁷⁵ Rescorla 2017: S5.1

elements. Fortunately, there is such a method. When it comes to natural language processing, we may simply assign different syntactic types to the context-sensitive elements of natural language. For example, when Jane says “I will do some reading” the “I” in her utterance may be given a unique symbol such as I_{Jane} . So, we can make use of alternative, non-indexical symbols in computation to replace indexicals that we encounter in natural language. (A similar story might be told about homonyms and the several references of a given name).

A second possible means of replacing indexicals in thought is given by Ball (2015). The proposal here is, roughly, that for locations, people or whatever it is that we initially wanted to treat indexically, we can have a constantly updating list of names. The most recent of which is then read off the top of the list. So, when it comes to “now” ‘A computer could be designed so that just after each signal, it creates a new file, labelled with the value from the counter. Information about the present time could be stored in the highest-valued file; information about past times could be stored in relevant lower-valued files.’³⁷⁶ So when I want to think something about the present time, I take the name with the highest value, when I want to think about “then” (in the past) a lower value name is selected. Analogous processes can be performed for other thoughts that we might otherwise think of as indexical. Each new file is not an indexical, it will always refer to the same time. Whether it is selected as “now” depends on whether it has the highest value or not.

As noted at the outset, the difficulty we face is that Claims (i)-(iii) seem incompatible with each other: claim (iii) seems incompatible with claim (ii) and the conjunction of claims (i) and (ii) means that the indexicals cannot just be removed. There are, it seems, three possible ways out. One could deny (i) or (ii), which would make the CTM and the *De se attitudes* compatible with one another. This is because the point which makes the two incompatible is that one is committed to indexicals and the other is committed to not having them. If we deny (i) or (ii) then we do not need any indexicals. In this chapter, however, I will defend (i) and (ii) and deny that (iii) causes a problem. The issue of central

³⁷⁶ Ball 2015: 356

importance to this argument is whether we can have indexicals on the CTM. I will argue that it is possible to have a CTM and indexicals, despite a tension between these views.

The challenge in denying (iii) seems to be to explain how computations could be sensitive to the differences in content between token expressions of a given type. That is, how a processor which reacts only to the type of symbol that it encounters could also be sensitive to the features that the token has in that context of use. For example, “I” is a type. This type has a character. That character is a function from context of use to content.³⁷⁷ So, with different contexts of use, the character can result in a different content. The struggle is to have this difference in content whilst the symbol remains the same. To replace the symbol with a non-indexical symbol is to deny that *de se attitudes* have an essential indexical component.

An indexical only has its content as a token, rather than as a type. Given that the CTM is a functional theory of mind, and the function of a symbol is usually determined only by the properties of the symbols considered as a type, how can the symbol *qua* token be considered by the computations? This is the challenge that I will take up in section 3.5. I will argue that one can accept the CTM without being committed to the mind being computational through and through. Instead one can make use of a weaker commitment, so that only some elements of the mind are computational. I will deny that indexical symbols need to be treated computationally on the CTM. First, I will defend claim (iii).

3.4.1 Defence of claim (iii)

Claim (iii) is the claim that computations are only sensitive to the syntactic *type* of a symbol; they are not sensitive to other properties that the tokens that they operate on may have. Here is an argument taken from an unpublished paper by Ball. It isn’t clear that he fully endorses the argument and his published work seems to show a more concessionary position according to which whatever can be done by an indexical account of attitudes can be done by a non-indexical account.³⁷⁸ However, the

³⁷⁷ Kaplan 1989: 507

³⁷⁸ Ball 2015

argument in his unpublished work presents a serious worry for the triad of views that I wish to endorse, so I will take it seriously here. I will refer to this as Argument 1. Argument 1 draws on the syntactic nature of the CTM to argue that the CTM is incompatible with indexicality.

Argument 1 can be put as follows.

1. If there are constituents of the language of thought that have an indexical character, then the computational theory of mind is not true.
2. The computational theory of mind is true.
3. Therefore, there are no constituents of the language of thought that have an indexical character (1, 2, MT)

It seems to follow from Argument 1 that the CTM is incompatible with indexicals in thought, so (i) – (iii) are incompatible. Kuczynski (2007) makes a similar point and argues that this means we should dispense with the CTM. As my aim is to show that there can be indexicals in thought, I will focus on Argument 1 as it provides a better foil to this position.³⁷⁹

The idea is that computations must treat all instances of a syntactic type in the same way. So, any tokens of that type which require different treatment are not possible. The key premise to defend is (1). This claim is plausible in so far as the Type Sensitivity Constraint (TSC) is true (see below). The TSC captures what is important about (iii) in a more rigorous way. The TSC follows from the CTM's reliance on mechanical processors to perform operations. A processor reacts in a mechanical way based on the shape of the symbol it encounters. It may copy or delete a symbol or type a new one, etc, depending on the symbol and its program. A second support for the TSC comes from the claim that if it were false then we would be liable to produce errors that we are not, in fact, liable to commit. These are errors of equivocation. I will elaborate on this below. I will first give a definition of the TSC and explain why it creates a tension with (ii).

³⁷⁹ Kuczynski 2007: 237

Argument 1 focuses on the TSC. It seems that if the TSC is true then we get the truth of premise 1 as well. The TSC is that:

‘There is some way of typing mental representations such that (i) if two representations differ as regards their content, then they are of different types; and (ii) computational mechanisms such as those postulated to account for reasoning are sensitive to type identities and difference among representation tokens’.³⁸⁰

For a given symbol type T, its tokens can express a different content each time, or not.³⁸¹ If not, then there is less motivation to think that T is an indexical symbol. This claim gets additional plausibility when we consider that indexicals do not refer as types.³⁸²

If the tokens do express a different content, then the TSC is violated. Condition (iii) is not met as the CTM entails the TSC. It seems that if the TSC is violated, then we are no longer using a CTM.³⁸³ This seems to require us to choose between (ii) and the CTM. Ultimately, I will claim that this is not so, but I will first present some reasons for endorsing the TSC.

One way to defend the TSC is to look at the nature of computation in the CTM.³⁸⁴ A CTM will operate according to computations. ‘A computation... is a formal operation on syntactically structured representations. Accordingly, a mental process, qua computation, is a formal operation on syntactically structured mental representations’.³⁸⁵ This understanding of computation makes these processes amendable to Turing-style computations.³⁸⁶ These computational processes are only sensitive to the syntactic *type* of the symbol that they are operating on (this is not true on all

³⁸⁰ Ball 2010: 5

³⁸¹ The possibility that they refer to different things need not be practical here, so it may be that the symbol always refers to one individual during that person’s lifetime. This does not mean that the symbol could not change its referent, just that it need not change its referent to still be encountering problems.

³⁸² Recanati 2012: 57

³⁸³ Ball 2010: 13

³⁸⁴ Ibid 4

³⁸⁵ Fodor 2000: 11

³⁸⁶ Borg 2004: 81

accounts of a CTM, but it is true on the E(CTM) which Fodor endorses, see chapter 2).³⁸⁷ This is because the computations use a “processor” which has set reactions to particular syntactic features.³⁸⁸ In which case, the processor will be unable to take other features into account. This includes features that the symbol might have, including any features that it has because it is the token that it is. (It is possible that there are different reactions to a given symbol, as one can change the “program” that the processor makes use of.³⁸⁹ I will ignore these complications as it seems the processor will still treat symbols as instances of a given type, rather than as a token).

If the processor could take non-local features of a token into account, such as the context in which it occurs, then we are no longer dealing with a computation that can plausibly be attributed to humans when they engage in certain kinds of reasoning. It does not seem plausible that human’s non-demonstrative reasoning processes, such as abduction, operate over a human’s entire belief system. However, if a given syntactic type could have its content changed depending on the context, then any number of beliefs that the agent has might prove to be relevant in determining that content. So, a process that determined the content of such a symbol would need to be able to check for *relevant* beliefs and to be able to treat the symbol accordingly in the future. This is not the sort of thing that a computational processor can do. Instead, the processor seems to be acting in a manner akin to an intentional agent who takes several features into account. At this point, one of the key advantages of the CTM is lost as we can no longer account for some rational inferences without already making use of rational agency. This seems to be a part of the rationale behind Argument 1.

So, it seems plausible that, on a CTM, if two symbols differ in content, then they are of different syntactic types. For one thing, we should expect that if two symbols have different content then this difference should be reflected in how those symbols are treated. For instance, the belief that there is

³⁸⁷ There are others who support this claim, for instance, Müller states that ‘At the syntactical level, the states or physical objects are taken to be tokens of a type (e.g. charge/no charge) and are manipulated according to algorithms’ (Müller 2008: 119).

³⁸⁸ Pinker 1994: 75

³⁸⁹ Cain 2016: 31. These programs can be understood as different algorithms being used on the symbol (Müller 2008: 119).

a postman at the door will (hopefully) elicit different reactions to the belief that there is a murderer at the door. So, these two contents will be treated differently. If a computational process is going to acknowledge this difference it will have to perform different operations over them. If computations are operations over different syntactic types, then one way for the process to acknowledge the different content is to make use of different syntactic types.

An alternative way for the difference to be accounted for is in the different tokens of a given type being treated by a different algorithm, but this seems unsatisfactory. Even if the tokens were treated by the same algorithm, we would still want them to have different results. It seems counter intuitive in dealing with logic and deductive inferences that we can have one symbol and treat it differently depending on what we want it to represent. It seems simpler to have more symbols. Given that computations are deducers, deductions and computations are comparable cases. Though, this point is not definitive as it does not make the use of several algorithms impossible, it just points out that postulating more algorithms makes things more complex than need be.

A further point in defence of the TSC is that if we were to deny it, then we would be liable to commit errors that we are not in fact liable to commit. For instance, 'I should be disposed to infer from *that is a cat* and *that is a dog* to *something is both a cat and a dog*'.³⁹⁰ The point is that if we could have a mental symbol whose reference varies with the context of use then the computations involved would not always (if at all) be able to track this change in reference. The processor would come across "that" and treat it as though it had the same content in each instance. In which case, we would get an error of equivocation as presented here. However, given that we do not equivocate as often as this theory would seem to predict, this is evidence against Indexicalism in a CTM.

This point is particular to certain sorts of indexicality, and it may not apply to those cases in which we have practical non-indexicality, as is the case with *de se thoughts*. This is because there would not be a prediction of equivocation. Given that the symbol for the self would always be in the same context

³⁹⁰ Ball 2010: 13

of use relative to what it refers to, i.e. it will always have the same content across all the contexts that it is likely to be used across, then it would not predict that there would be these errors in equivocation. (This is not to say that it is impossible for there to be cases in which the prediction would change. For instance, some sort of brain transplant might have this effect. However, in this case it is not obvious that the prediction of some sort of equivocation would be incorrect).

Given that the mental “I” that will appear in the person’s belief state is not practically indexical, it may seem like an unfair case to apply to Argument 1. However, Ball seems explicit in taking his considerations to apply to such thoughts as well. He says that ‘...the reason that the thought Perry expressed by ‘I am the messy shopper’ was cognitively distinct from the thought that he expressed by ‘John Perry is the messy shopper’ is because the two thoughts have different characters; in particular, the ‘I’ thought has a context-sensitive character. So, the ‘I’ thought is indexical. I will argue that this explanation cannot be correct if thinking is a matter of computation’.³⁹¹ This suggests that Argument 1 is taken to apply to indexicals in general, and therefore to “I” thoughts as well, when these thoughts are indexical. So, it seems we cannot have (ii) if we have (iii).

3.4.2 Responses to Argument 1

In this section I will consider some responses that might be made to Argument 1. One possible response is a modification the TSC. An alternative TSC might allow for indexicals and still be consistent with the CTM. The proposed change goes as follows:

Type Sensitivity Constraint* (TSC*) There is some way of typing mental representations such that (i) if two representations differ as regards their content, then they are of different non-indexical types, or of the same indexical type; and (ii) computational mechanisms such as those postulated to account for reasoning are sensitive to type identities and difference among representation tokens of non-indexical types, and

³⁹¹ Ibid: 4

sensitive to type identities and differences as well as other associated information among representation tokens of indexical types.³⁹²

Using TSC* might suggest that indexical symbols can have additional features added to them to allow the processor to take the context of use into account. One is to add an additional feature to each of these symbols, such as a subscript. One possibility, when considering *de nunc* cases, is to give a time stamp to each tokening of NOW to allow the processor to track it appropriately. ‘...for example, we could imagine representations prefixed with a sort of temporal quantifier or operator (AT 2:00, THE MEETING STARTS NOW), or representations that link the NOW to a time via a sort of parenthetical (THE MEETING STARTS NOW (2:00)).’³⁹³ This response allows the computation to take the context into account by associating relevant information with the token indexical.

This proposal is like adding a subscript to the logical form of the thought. So, in the case of *de se thoughts* an equivalent proposal would be to add a subscript that identifies the person in question. Rather than having a mental “I”, one might instead have a mental “I_{NN}” or a symbol that has a subscript that is appropriate for the individual. Here, the subscript types the referent of the symbol to allow it to get a content, rather than a changeable character and the context of use determining the reference.

However, we no longer have indexicals in this case. These are not types whose tokens could express different contents; they are now different types, each with a fixed content across contexts of use. Sticking to the time example, Ball says that ‘But now it looks like for the purposes of computation, we have two distinct types: the *nows* tokened at 2:00 are type-distinct from the *nows* tokened at 3:00....In general, any computational system that has the resources to avoid making the bad inferences in question will end up treating the alleged indexicals as of distinct types’.³⁹⁴ Each symbol is then a distinct type because they include the time to which they refer as a part of their symbol. In

³⁹² Ball 2010: 13

³⁹³ Ibid 14

³⁹⁴ Ibid 14

this case, we no longer have indexicals. It seems that a similar move can be made with other indexicals like “here” and “I”. In the “I” case, it seems we could have I@X, where X is the agent referred to. It seems that we now have a version of Millikan’s view, and have denied (ii) is true. So, the TSC* does not present an alternative to the TSC which also allows us to have genuine indexicals. In which case, it seems that we still have (1). So, the TSC* does not allow for genuine indexicals.

Another way in which one might respond to Argument 1 is to make use of computational memories. If the computer has a short-term memory it can use this to determine the reference of the relevant indexical. The worry for Argument 1, on this model, is that it can make sense for a computation to be sensitive to the symbol *qua* token. It may be a part of the character of the language of thought “I” that it refers to the self which is most recent in memory, or which has the highest context value, or whatever.³⁹⁵ So, when it comes to determining a single referent, the computer can search as follows:

...an initial search request which is parameterized by the focus threshold is evaluated.

Thus the search ignores entities which satisfy the constraints if they have context activations lower than the threshold, and also ignores entities which satisfy the threshold condition but not the constraints. If this initial search fails to locate any candidate referents, then a second search is made without the threshold condition. If more than one entity is located by either search, the context activations of the entities are compared and the one with the highest context activation is chosen.³⁹⁶

This is a little opaque, but an illustration may help. If one is trying to figure out who a given “he” refers to, one first needs to satisfy the constraints imposed by this. Here, a constraint would be that the referent is male. It then seems that the most salient individual who satisfies the constraint is selected to be the referent. Even if this story is not correct, it suggests that there are ways for

³⁹⁵ Alshawi 1987: 37

³⁹⁶ Ibid 37. Another alternative is presented in Lappin (2003: 108) according to which certain features can be recorded from a discourse to be used later in understanding the utterances of a dialogue. This would be useful in the case of anaphora.

computations to make use of memory and therefore for the computation to be sensitive to the symbol as a token. So, we may have a symbol in Mentalese whose function is to initiate a search for the relevant item. This may be a more promising way of spelling out the TSC* than making use of a subscript.

This may avoid the TSC's implication that there can be no indexicals in thought. Instead of that, we are requiring that the indexical draws on another non-indexical concept, which is determined contextually. In the case of *de se attitudes* the non-indexical concept would presumably be one that the agent has for themselves, not the mental "I". In which case, it seems that we have an indexical which can be used to pick out oneself as a referent.

Responding like this seems to appeal to a distinction between fundamental and non-fundamental levels of thought. If a thought is true in virtue of another thought, then it is not fundamental.³⁹⁷

Distinguishing thoughts like this does not help to resist Argument 1, however. The distinction helps Argument 1 in the following way. Suppose that the symbol which starts the search for a contextually salient and appropriate entity is a non-fundamental symbol. It seems that the symbol that is found because of this search will have to be fundamental. If it is not, then there will have to be another search for an appropriate signal. If that next symbol is again an indexical then we are in danger of being launched on a regress in thought.³⁹⁸ If this regress is motivating Argument 1 then we can make Argument 1's conclusion into a weaker, but more plausible claim. This is that there can be no indexicals at the fundamental level of thought, so something like the TSC must be true of this level. If the TSC were false at this level, then we might be at risk of this regress. This ought to make the TSC more plausible. It also means that this version of the TSC* does not apply at the foundational level of thought. In which case there is no character-underdeterminacy at the fundamental level of thought

³⁹⁷ Evans 198: 112. See also chapter 1.

³⁹⁸ This can be understood as an instance of the regress argument outlined in chapter 2.

and Argument 1 would still serve to undermine the idea that there cannot be context-sensitivity in thought.

To conclude this section, an argument can be made to the effect that if the CTM is true, then there cannot be mental indexicals at a fundamental level of thought. This is on the basis that a CTM uses a processor that can only react to symbols by taking their syntactic type into account. However, an indexical view would require taking the symbol *qua* token into account.³⁹⁹ This requirement seems to suggest that the CTM cannot allow for mental indexicals. In this case, claim (ii) seems to be false, and that threatens claim (i). In the following section I will attempt to show that this threat to (ii) and (i) can be avoided without giving up on the CTM.

3.5 Rejecting Premise (1)

I will deny that premise (1) is true. This premise stated that if there are constituents of the language of thought that have a context-sensitive character (that make them indexicals), then the CTM is not true. I will argue that the CTM can allow for non-computational processes to occur outside of mental modules. This is the version of the CTM that is defended by Fodor (1981, 2001). On this version of the CTM one can allow that computations are not able to treat symbols as anything other than instances of a given type without taking this to conflict with the *de se indexicalism*. I will not claim that there is a module that deals specifically with *de se attitudes* which then must give an output to a general reasoning centre to get content. Rather, I want to claim that the problem Argument 1 raises would only apply to those *de se thoughts* when they are processed within a module. This might happen when someone wishes to express an “I” thought, and this must go through a speech

³⁹⁹ A possible line of argument would be that it seems to be strange to have evolved a computational system which has an indexical element that does not change its referent. If this were the case, why not have a non-indexical element that would not change its reference, in which case we have something that seems to be computationally simpler and does not have any unnecessary complexity (having only a content, rather than a character and a content). My response to this sort of worry will, I think, be the same as the one that I will argue for in section 3.5. This is to reject the idea that the mind is computational all the way through so that what is computationally simplest is not the deciding factor in our understanding of the mind.

processing module. These instances need not be problematic. That the module is not able to treat symbols *qua* token does not matter, so long as a general reasoning capacity can do so.

However, this also is not enough. It cannot be the case that the non-modular part of the mind saturates the indexical and removes this context-sensitive part. To do so would be to deny (i) or (ii). In which case, the solution will not have been successful. If we take saturation to be the process by which indexicals are given a content, then an account of how this thought can refer without being replaced will be required. If the indexical is saturated, and replaced with a non-indexical element, then we will have lost the indexical element. In which case, we are no longer dealing with a fundamental *de se thought*. In the next section I will discuss ways of constraining the CTM to show that Premise 1 is not true.

3.5.1 Restricting the Scope of the CTM

I will argue that the CTM can accommodate cases of indexicals without requiring that it processes their content in a computational manner. Instead, the symbol can be moved out of the module (if it was ever there at all) to be dealt with in a non-computational manner. This move requires endorsing a modular theory of mind according to which we have domain specific modules (i.e. modules that are dedicated to a particular task and a particular kind of information) and a general information processing part of the mind that can deal with other processes that are not dealt with by modules.⁴⁰⁰ This general part of the mind will be able to draw on relevant, available, information to whatever its task is at that time. I will first give some motivation for this understanding of the CTM. I will then go on to deny that (1) is true.

⁴⁰⁰ One concern with this problem is that it is possible that sentences like “Jill is happy” be processed by a module to produce a content. On the view considered here the same does not seem possible for sentences like “I am happy” (were the person comprehending this is the one that uttered it, otherwise I do not think that there is a difference between these two sentences in terms of their possible comprehension by a hearer). These two sentences are very similar to one another, so it is not obvious it should be possible to get truth conditions from one instead of the other. In response to this concern I would like to claim that the two sentences are not comparable for the same reasons that *de se attitudes* are considered unique. So, we should not expect them to behave in the same way as other sentences.

The rationale for this limited CTM stems from the Frame Problem.⁴⁰¹ This is important because the limited CTM should not be motivated only by Argument 1. That would make the solution seem *ad hoc*. To briefly recap S2.1.1.2.1, the Frame Problem begins from the idea that when we are making some, typically non-demonstrative, inferences to a belief we try to make use of premises that are relevant to that task. Here, a relevant proposition can be understood as one that ‘if it *were* attended to *would* affect the estimated subjective probability of the belief [derived by this reasoning]’.⁴⁰² The problem is that this property is not one that can be had in virtue of the syntactic properties of a belief or thought.⁴⁰³ This means that computations cannot be responsible for selecting the relevant premises for these sorts of inferences and tasks. So, computations cannot be responsible for these inferences. This supports the case for there being a general, non-computational domain of the mind. We can, however, accept computational processes that do not have to deal with the Frame Problem. These will have to be performed over areas where the range of possible information is already constrained, so that the selection of relevant information is not done as a result of a computation.⁴⁰⁴ Instead, ‘...the information that is available to perform a task depends on which task it is; and the constraints in virtue of which this is so are “architectural”’.⁴⁰⁵ The result is that ‘modules are informationally encapsulated by definition’ and so they do not have to select the information that they make use of.⁴⁰⁶ So, we can preserve a CTM without having to claim that the mind is thoroughly computational (see Chapter Two for more discussion on the Frame Problem and modules).⁴⁰⁷

Now that a modular CTM is on the table, it is possible to give a way of denying (1) and therefore deny that (iii) causes a problem. The solution here is like one that Borg (2012) makes use of when faced

⁴⁰¹ At least, it stems from Fodor’s version of the Frame problem; see Kamermans and Schmitts 2004 for other variations of the Frame problem.

⁴⁰² Fodor 2010: 116

⁴⁰³ Ibid 124

⁴⁰⁴ Fodor 2000: 64

⁴⁰⁵ Ibid 63

⁴⁰⁶ Ibid 63

⁴⁰⁷ Alternatively, one might want to dispose of the CTM altogether. In this case, we also do not have an inconsistent triad as (iii) ceases to be relevant. Though, in this case one needs to propose an alternative theory of mind which can accommodate indexicals.

with a similar problem. '[Semantic Minimalism] wants it to be the case that the processes involved in recovering the semantic content of any well-formed sentence are deductive, computational processes – processes which look to the formal properties of representations and not to their content'.⁴⁰⁸ However, demonstratives often seem to require some further input from the context, such as the intention of the person using it, or what is contextually salient at this time, in order to select their referent. Getting this information requires appealing to the context. This seems to be a similar problem to the one faced here. A content is not accessible via the computation alone, because getting a content seems to require information and processes that are not available to the computation. In our case, it is the self that that token refers to, rather than interlocutor's intentions doing the work. (Borg's problem differs in that she needs a computation to yield a content in accordance with Semantic Minimalism, so that there must be a content making minimal use of the (objective) context. There is no such restriction here. A second difference is that Borg's problem arises when processing input from another speaker, whilst the problem here arises whenever one is thinking about oneself in a *de se* way, in a module. A third difference is that it does not seem to be as much reason to think that there is a dedicated module for *de se attitudes*).

Borg's solution to this problem is to allow for the character to stay fixed and allow for a weak linguistic content, with richer details to be determined outside of a (linguistic) module. So, when considering demonstratives, the idea is that there is a content to the demonstrative that can be turned into '*the actual object referred to by the speaker with this token of "that"*'.⁴⁰⁹ This is a linguistic content and on Semantic Minimalism it is enough to be truth-conditional. Producing this linguistic content is all that the linguistic module must do. Finding a non-linguistic way of presenting the actual object is a task for some other part of the mind, most likely the non-modular general processing part of the mind.⁴¹⁰

⁴⁰⁸ Borg 2012: 134

⁴⁰⁹ Ibid 135

⁴¹⁰ Borg 2004: 199

With Argument 1, a similar response might work. Whilst a module that the mental “I” might appear in is not able to take this symbol into account *qua* token, it can pass this on to the non-modular part of the mind and here the content of that symbol can be worked out, ‘...that is to say, even at the level of thought we need to distinguish content and character’.⁴¹¹

This response seems to require that we have a module dedicated to having *de se attitudes*, from which the indexical is shipped out.⁴¹² I do not want to be committed to this claim. Rather, I want to concede that there could be a problem for indexicals on a CTM. But I want to say that we only need computation within modules. So, the problem that Argument 1 raises only occurs within modules. So, if we grant that there may be cases in which this is a problem for a modular CTM, we still need not be too concerned. If there are cases in which the *de se thought* occurs within a module, its content can be worked out in a general processing system. So, there is no inconsistency.

There is a worry that this approach does not do enough by way of response to Argument 1. I have not shown that the TSC is incorrect, only that it does not apply to all the processes that happen in the mind. It bears emphasising, however, that this still shows a gap between Argument 1 and the conclusion that there cannot be indexicals in thought. Also, this understanding of the CTM is not novel and is well known from Jerry Fodor who is one of the chief proponents of the CTM. In his words ‘...when I wrote books about what a fine thing CTM is, I generally made it a point to include a section saying that I don’t suppose that it could comprise more than a fragment of a full and satisfactory cognitive psychology’.⁴¹³ A similar claim is that ‘Even if input systems are domain specific, there must be some cognitive mechanisms that are not’.⁴¹⁴ From this we get the idea that the CTM need not be expected to be exhaustive of our understanding of cognition.

⁴¹¹ Borg 2012: 135

⁴¹² I would like to thank Professor Borg for highlighting this point for me.

⁴¹³ Fodor 2001: 1

⁴¹⁴ Fodor 1983: 101

Furthermore, it may be that we should not expect *de se attitudes* to be properly instantiated by modules in the first place. *De se attitudes* seem important because we use them for practical deliberation, introspection, self-criticism, self-knowledge, determining what our place in the world is and so on.⁴¹⁵ For any one of these tasks we may require access to a range of considerations which do not seem compatible with a modular account of the mind. This is owing to the range of information being non-domain specific. It can include information relating to one's environment, one's past, one's relations, what one would like the world to be like and so on. This variety of information makes a general processor seem like the most likely home for *de se attitudes*, as their use can draw on and be related to a great deal of information. This central position of *de se attitudes* suggests that we should not attempt to understand them as operating in a computational manner. Their use in relation to a wide variety of further beliefs suggests that they cannot but be subject to the Frame Problem.

This way of responding to Argument 1 can also avoid predicting errors that we are not, in fact, liable to commit. Recall the worry that indexical symbols in thought would predict that we are liable to errors in equivocating that we do not, in fact, commit. For instance, "That is cat, that is a dog, therefore that is a cat and a dog". This prediction can be avoided by avoiding a commitment to a thorough CTM. If the process is not computational then it is not a requirement that all instances of a given type be treated as though they had the same content. So, each instance of "that" can be treated as a token with its own content. So, this account need not predict errors of equivocation.

We also do not need to deny the TSC. We need only restrict its scope to those elements that occur within whichever modules it is that we make use of and therefore constrain it to the computational parts of the mind. This move also opens the possibility of there being indexical *de se attitudes* at the fundamental level of thought.

⁴¹⁵ De Gaynesford 2006: 173-4, Guillot 2016: 138

3.5.2 Summary

In sum, this solution requires a modular account of mind on which some of the mind functions non-computationally. If the mind were computational through and through, then passing the indexical symbol on to another part of the mind would not be helpful. There would continue to be the problem from Argument 1. Instead, there needs to be some component of the mind that does not function in computational way. The Frame Problem also gives us reason to think that there are such parts of the mind as this, so this is not an *ad hoc* move. From here, we can allow that the richer parts of the context can be taken into consideration in determining the content of a thought. So, Argument 1 only has an impact when the *de se* thoughts appear in modular parts of the mind, if they appear there at all.

Here, I want to concede that computations do need to meet the TSC but deny that a CTM needs all its processes to occur computationally. Instead, only the modules operate computationally. So, the idea that we might have thoughts with a character need not be incompatible with the CTM. (It may be that all a computational module can do here is provide a semantic constraint, which constrains which propositions could be expressed by a term without determining a proposition).⁴¹⁶

One response that might be made here goes as follows. I have presented a dilemma between the CTM and *de se indexicalism*. To resolve this, I have opted to dismiss the CTM in favour of *de se indexicalism*. But, why not go the other way and deny *de se indexicalism* and keep the CTM? There are a couple of things to say in response. One is that I do not think that this is what I have done. I have not argued that the CTM should be disposed of. I have just argued that the scope of computations should be restricted to modules, and that the mind is not entirely modular. Of course, one might say that this makes little difference. Why not keep a thoroughly CTM and lose *de se indexicalism*? One reason is that the CTM must deal with the Frame Problem, and it does not seem able to if it is thoroughly computational. Another is that one must show that the problem that Ninan

⁴¹⁶ see Harris (2018: 6) for discussion.

identified stemming from *de se attitudes* is not unique to *de se attitudes* to undermine the case for *de se essentialism*. But this does not seem to be an easy thing to do, as I have argued above. If similar examples occur for other indexicals this may also show that the class of essential indexical is broader than just those in *de se attitudes*.

One objection that might be made here is that I have not offered a solution to Argument 1 as the general reasoning processes are difficult to describe or understand in any detail, whereas it is an advantage of the CTM that its processes can be described. In response to this, I would like to emphasize that the move to general processing is motivated by the Frame Problem. Whilst there is a price to pay in terms of which processes we can understand computationally, this is a price that we are already paying. Whilst this means that we do not fully understand how *de se attitudes* should work, it seems that if we adopt (i) – (iii) then however they work they cannot work computationally. Furthermore, my aim was to show how it is possible to maintain a CTM whilst also allowing for indexical thoughts in the form of *de se attitudes*. I have presented a means of doing that. This view differs from that introduced by Argument 1 as indexicals are a possibility on this view.

3.6 Conclusion

To conclude, I have argued that there is not a significant tension between the CTM and *de se indexicalism* if one endorses a modular CTM. This is because the CTM seemed to make it impossible for there to be indexicals at the level of thought, as computations are only sensitive to a symbol's syntactic type, rather than to the symbol *qua* token. But being able to consider the symbol *qua* token seems to be required by *de se indexicalism*. However, if the computation does not have to deal with indexical mental files as tokens, then there does not seem to be a problem. The computation can instead pass the symbol onto the non-computational part of the mind, if it is there in the first place, which can then take external factors into consideration to yield the content for that token. In this case, there is not a contradiction between the CTM and *de se indexicalism*.

This removes a barrier to accepting (i) and (ii). What's more, if the defence of these points is convincing, then it seems that we should accept them. This means that there can be context-sensitivity at the level of thought. Additionally, these indexical thoughts cannot be saturated away, without a loss in the attitudes explanatory power. This makes those thoughts both fundamental and character-underdetermined.

4. Searles Background Theory of Intentionality: Cause for Concern

4.1 Introduction

In the previous chapter I discussed Character–Underdeterminacy and argued that the case of *de se attitudes* makes a compelling case for context-sensitivity at the level of thought. This appears to be so even on a computational theory of mind (CTM), provided that the mind is not thought of as being thoroughly computational i.e. where computation is restricted to modules, which feed into a non-computational processing system. What’s more, these *de se* thoughts did not seem to be true in virtue of an additional, more foundational, thought. This provides core evidence for my central claim, that it is possible to have foundational sentences in the language of thought that are context-sensitive.

In this chapter, I will investigate another kind of underdeterminacy and consider whether it is possible to make sense of this kind of context-sensitivity applying to foundational thoughts. I will be concerned with type-underdeterminacy, where this is understood as:

Type-Underdeterminacy: A non-indexical structured representational item S is type-underdetermined if and only if there are tokens of S that have distinct truth-values.⁴¹⁷

An example of type-underdeterminacy is the natural language sentence “Sam is a saint” which may express either the claim that *Sam is a saint* in the sense of having been canonised or *Sam is a saint** meaning that Sam is kind and friendly. In which case, for the same circumstance of evaluation the sentence can convey different contents and thereby change the truth-value.

To consider this version of underdeterminacy I will make use of Searle’s understanding of Really Radical Contextualism (RRC). On this view, the content of an intentional state, e.g. a tokening of a given thought, is dependent on a set of Background of assumptions (when discussing what Searle calls the Background I will use a capital “B” to show that this is distinct from other, non-technical

⁴¹⁷ Jaque 2017: 3. This has also been described as context-sensitivity by MacFarlane (2009: 232). I will take tense to be an indexical, and so not a proof of type-underdeterminacy.

uses of the word “background”).⁴¹⁸ This makes Searle’s position, which I will refer to as the Background Theory, a natural test run for type-underdeterminacy as Searle explicitly endorses it. Whilst the failure of this view would not rule-out the possibility of type-underdeterminacy at the level of thought, it would be damaging for the claim that there is type-underdeterminacy in thought since Searle’s account provides what I take to be the most plausible version of this sort of view. On the other hand, if the Background Theory is successful in avoiding significant challenges (such as the regress challenge to be explored below), then this would be very good reason to endorse the claim that context-sensitive foundational thoughts are possible.

In this chapter I will outline Searle’s position and his arguments. I will then introduce an idea from Perry who distinguishes between what he calls an aboutness relation versus a concerning relation (these ideas will be spelt out in 4.3) and explore how this might be used to avoid a regress. As the concerning relation does not require additional features for the thought to be about a state of affairs, I will suggest it shows that an additional thought need not be generated and thus that the challenge of regress levelled against context-sensitive thoughts can be avoided. Next, I will turn to consider some objections to this account of type-underdeterminacy. For instance, I will consider whether this account is committed to a problematic kind of holism. Another issue is that it is unclear whether we should have something other than thoughts as the basis for intentionality. I will conclude that this promising account of type-underdeterminacy in thought is successful in avoiding a regress but at the cost of having to lower the standard at which we would normally count as agreeing and disagreeing with one another and with oneself over time. This cost motivated by considering that communication does not require that two people share identical contents.

4.2 Searle’s Background Theory

In this section I will outline Searle’s Background Theory and his arguments for it. I will argue that whilst his arguments have some shortcomings, these can be adjusted for. This will motivate Searle’s

⁴¹⁸ Searle 1980: 231

conclusion that any intentional state underdetermines when it would be true (thus it underdetermines its content). Whilst I do not need to endorse this strong conclusion to prove my central claim (that some foundational sentences in thought *could* be context-sensitive) assuming such a pervasive account of type-underdeterminacy is useful for my approach as there are no underdeterminacy free thoughts to appeal to on such a picture. In which case, there is no foundational thought free of context-sensitivity to appeal to. If a regress can be avoided on this view, then it follows that foundational thoughts can be context-sensitive. I do not take myself to have proven that Searle's view is correct, rather what I will try to show is that a regress can be avoided on such a view.

On the Background view, an utterance of a given sentence can only have determinate truth-conditions if it is understood within a Background of other states which can themselves be either intentional or non-intentional. That is, the Background can consist of other things which we can represent but do not need to. Searle claims that:

...as members of our culture we bring to bear on the literal utterance and understanding of a sentence a whole [B]ackground of information about how nature works and how our culture works. A [B]ackground of practices, institutions, facts of nature, regularities, and ways of doing things are assumed by speakers and hearers when one of these sentences is uttered or understood.⁴¹⁹

Here, the focus is on natural language, but it gives one an idea of what the Background is. It is a set of assumptions that one makes about the world and context which we inhabit. It is also important that the Background includes a very broad range of assumptions that one might have about a given situation and we need not be entertaining these assumptions consciously.⁴²⁰ This means that there is

⁴¹⁹ Ibid 226-7

⁴²⁰ Ibid 228

a very low bar to entry to the Background. Any assumption that one might make about the world could form a part of the Background whether one realises it or not.

This does not mean that the Background includes all the assumptions that a person could make about the world, which would be troubling as that is an unfeasibly large set of assumptions for an individual to have. Instead the Background consists of assumptions that an agent may act on, even without realising that they are doing so. For instance, the assumption that grass does not eat people is probably in most adults Background assumptions even if they never realise that it is. The Background does not have to be internally consistent either.

The Background seems indefinite and not composed only out of representations, though any assumption in the Background can be represented.⁴²¹ This means that the Background can be an incredibly broad range of assumptions. It is, therefore, difficult to delineate what is a part of a person's Background or not. It does not have to consist of assumptions that are consistent with one another or that ever consciously entertained. It may also include assumptions that are easily inferred from other assumptions that one has.

To distinguish which assumptions are a part of someone's Background from those that are not, I propose to use a criterion based on manifestness. An assumption is '...manifest to an individual at a given time if and only if they are capable at that time of representing it mentally and accepting its representation as true or probably true'.⁴²² This capability is not so broad as to include anything that could be represented by anyone. It depends on what A can represent in that circumstance.

Determining exactly what does and does not get into the Background is a difficult problem that I will return to in section 4.4.1.

There are a few other features of manifestness that are worth noting. One is that manifestness does not require that the assumption be represented. It just needs to be possible that the agent

⁴²¹ Ibid 228

⁴²² Carston 2002: 67-68

represents it. Another feature is that manifestness is a matter of degree. It may be more or less difficult for an individual to represent an assumption at a given time and they may be more or less likely to accept it as probably true. These features make it well suited to playing a role in making sense of the Background.⁴²³ In which case, we can say that an assumption P is a part of agent A's Background assumptions at time t if and only if it is possible for P to be manifest to A at time t.⁴²⁴ This is what is possible relative to the context in which A finds herself. This provides a way of distinguishing assumptions that are not a part of someone's Background from those that are, whilst maintaining that it is very easy for an assumption to be a part of someone's Background.

Now that I have said a little about what the Background is, I will give an account of Searle's view. When discussing natural language sentences, Searle puts his view as follows 'the literal meaning of a sentence only determines a set of truth conditions given a set of [B]ackground practices and assumptions'.⁴²⁵ Here, Searle commits himself to the type-underdeterminacy of natural language sentences. He also claims that the Background is a necessary condition for removing type-underdeterminacy.

Searle goes on to claim that 'Given different [B]ackgrounds, one and the same sentence with one and the same literal meaning may determine different conditions of satisfaction, and given some [B]ackgrounds a sentence or concept may not determine a definite set of conditions of satisfaction at all'.⁴²⁶ The "conditions of satisfaction" are the conditions under which a given speech act would have satisfied a set of conditions based on its illocutionary force. In the case of an order, it is satisfied when the order is carried out as intended. The conditions of satisfaction for an assertion are its truth-conditions. I will continue to focus on thoughts that can be true or false. Given that the assumptions that compose a person's Background could remain fixed whilst the content of a sentence type may

⁴²³ Ibid 68

⁴²⁴ For some purposes it will be best to take it to be a short interval of time rather than as single moment, as it may take one a second or so to form the thought.

⁴²⁵ Searle 1980: 227

⁴²⁶ Ibid 231

vary, it seems plausible that an assumption makes more of a difference to the content of a sentence the more manifest the assumption is.⁴²⁷ That is, the assumptions in the Background might remain fixed whilst the degree to which they are manifest varies. The more manifest an assumption is, the more of an effect it will have on determining the content of a representation. For example, one might have the thought I WANT A PET DOG with two different Background assumptions. The first assumption may be that Labradors make good pets. The second is that Whippets make good pets. If one of these assumptions is more manifest than the other, that may affect the kind of pet dog one takes oneself to want.

Importantly, the Background does not function as an unarticulated constituent to the thought or to modulate the thought.⁴²⁸ The Background does not function as a means of adding more components to a representation or altering ones that were already there. It does not add more character to a representation to determine a content. It instead acts to determine a content in some other way. I will suggest that the concerning relation is a good way of understanding how this works.

These quotes show that Searle endorses claim (a) of RRC, that all (or almost all) natural language sentences are context-sensitive. He also maintains (b), that all (or almost all) thoughts are context-sensitive. Searle thinks that ‘...the features we have cited [the dependence of content on a Background] are features not just of semantic contents but of representations generally, in particular they are features of intentional states, and since [content] is always derived from intentionality, contextual dependency is ineliminable’.⁴²⁹ As thoughts are a kind of representation, it follows that Searle also endorses (b).

⁴²⁷ Carston 2002: 66

⁴²⁸ Searle 1980: 228

⁴²⁹ Ibid 231. See also Searle (1979: 216) for similar claims. Searle’s claim is phrased in terms of meaning always being a derived form of intentionality. I have changed it to content as that seems to be Searle’s claim as it would appear in my terms. If it were meaning in my terms then Searle claim would be more difficult to make sense of, given that he seems to take meaning to remain fixed whilst content varies. I will continue this practice when quoting Searle to show my interpretation of him. Searle (180: 226) claims we learn words like “cut” ‘...with their *common* meaning determining *different* truth-conditions’ which supports my reading him this way.

Having sketched Searle's view, I will now give his argument for it. His argument begins with several sentences, each with the word "cuts" in them. After excluding non-literal uses of "cut" Searle considers the following sentences "Bill cut the grass", "the barber cut Tom's hair" and "Sally cut the cake".⁴³⁰ In each case, "cut" seems to describe a different action. The first would involve a lawn mower, the second some scissors and the third a knife. Searle notes that '...in those sentences one and the same semantic [meaning], expressed by the word "cut", occurs in each sentence; and yet it seems to make a different contribution to the truth condition of the sentence in each case'.⁴³¹

That the contribution to the content can vary is something that Searle takes to undermine the compositionality principle. The version of the compositionality principle that he has in mind is the composition of content, according to which the content of a representation is determined by its parts and the way in which they are combined. Searle's point appears to be that the content of a word, "cuts", is not determined by that word's character (or content, if it has any on this view) as the word remains constant, and so its character does as well, whilst the utterance's content changes.

Searle also considers some more extreme cases, such as the sentence "Snow is white". Whilst we would not ordinarily take this sentence to be context-sensitive Searle argues that this is just because the assumptions that are usually relevant for determining its content rarely, if ever, change. To show how a change in assumptions might cause us to reconsider the content, Searle considers a case in which the earth is showered by radiation that changes both the wave length of light that snow reflects and the human visual apparatus, so that the snow does not appear to change colour. Nevertheless, research finds that snow reflects a wavelength of light that would have been called "chartreuse" prior to the radiation.⁴³² This case suggests that there are two possible contents for "Snow is white" to convey. For one content the thought that SNOW IS WHITE is false as the content it conveys is something like *snow reflects light that appeared white to humans prior to being radiated*.

⁴³⁰ Searle 1980: 221. This is not an exhaustive list of the sentences that Searle considers but it is sufficient here.

⁴³¹ Ibid 223

⁴³² Ibid 230

For the other content it is true, if the content is something like *snow now appears to humans to be white*. Which of the two contents is conveyed by the utterance is not forced on us by the sentence itself. This suggests that the content of this sentence is not simply the result of composition.

The rest of Searle's argument focuses on ruling out alternative explanations for the change in content that these utterances seem to undergo. If he is correct that there are no alternative explanations, then it seems that the parts do not determine the content of the whole and we have type-underdeterminacy. It would show that the sentence type can have varying truth-conditions and therefore values. He considers ambiguity, vagueness and whether the content of "cut" varies because of sentences that it is a part of.

Before discussing alternative explanations, it is worth reiterating that Searle takes these considerations to apply not only to natural languages sentences but to beliefs as well. Searle takes it that the arguments for the context-sensitivity of natural language sentences 'are also arguments for the context-sensitivity of the *beliefs* that Bill cut the grass... The content of those beliefs determines the conditions of satisfaction that they do only against a [B]ackground'.⁴³³ Thus Searle endorses the claim that all thoughts are context-sensitive as well.⁴³⁴

The first explanation that Searle aims to rule out is ambiguity. "Cut" is not ambiguous because 'in each of its occurrences it involves a common semantic [meaning] roughly involving the notion of a physical separation by means of the pressure of some more or less sharp instrument'.⁴³⁵ This is in contrast to well-known ambiguous words like "bank" which can have two distinct meanings, such as financial institution or land next to a river. The best way to understand Searle's claim is that there is a common character in each case, rather than two distinct characters as we find in the case of "bank".

⁴³³ Ibid 231

⁴³⁴ This is stronger than claim (b), see chapter 1, that all or nearly all thoughts are context-sensitive, but Searle's position is similar enough to call him a really radical contextualist.

⁴³⁵ Ibid 224

Another alternative that Searle considers is that ‘these different interpretations are determined by the different arguments – grass, hair, cake...’.⁴³⁶ In which case, the content of the sentence is still composed by the meanings of its parts, it is just that there are other parts of the composition that change the content that “cut” has. This is a plausible alternative as each of the examples that Searle has given are of a different composition. Searle denies that this is what is happening ‘...for it is easy to imagine circumstances in which “cut” in “cut the grass” would have the same interpretation it has in “cut the cake”, even though none of the semantic contents of the words has changed’.⁴³⁷ In which case it is possible that the sentence remains constant whilst the content varies. One may need some slices of earth with grass on, for some reason.

There are numerous reasons why an opponent of Searle would not find this compelling. One is that the ability of a sentence to communicate different contents under different conditions does not show that there is not a common content to each of these utterances that is derived from the sentence according to its composition (see *Semantic Minimalists*, chapter 1). This common understanding might be something like the definition Searle considers when ruling out the possibility of “cut” being ambiguous, namely, that to cut is to separate using a more or less sharp instrument. In which case, this might be a content that remains a constant part of the sentences that “cut” appears in. To rule this out Searle would need to assume that the semantic content we pay the most attention to, or are most aware of, is the only one that that sentence has. This assumption would require defence to convince a Minimalist.

Searle considers a Minimalist response when he says that “cut” is not simply vague. This alternative explanation would make use of a contrast between speaker meaning and content. He claims that ‘According to this account the examples are standard cases of the difference between less precise literal sentence [contents] and more precise speaker’s utterance [contents]’.⁴³⁸ The idea would be

⁴³⁶ Ibid 224

⁴³⁷ Ibid 224

⁴³⁸ Ibid 225

that the word “cut” is vague and always makes the same contribution to the content of the sentence. In which case the content still composes. The different contents that we pick up on would be the speaker meanings. Searle aims to rule this possibility by arguing that there is no minimal content conveyed here.

For Searle, this explanation predicts that other such sentences should be comprehensible, but they are not. Searle considers the sentence “Max cut the sun”.⁴³⁹ Searle argues that if there was a minimal content to “cut” then this sentence would make sense. He does not think that this sentence conveys any comprehensible content as he ‘does not know what truth conditions are supposed to be determined by these examples’.⁴⁴⁰ Given that Searle is a competent speaker of English and understands each of the words, that he does not understand the content of the sentence suggests that the content of the sentence is not determined only by the parts of the sentence and the way in which they are composed. Searle takes it that the lack of Background assumptions on what it would be to cut the Sun are to blame for his lack of comprehension here.

A response to this argument is that whilst it is hard to make sense of how someone would go about cutting the sun, it seems that there will be cases in which it is clearly false that Max has cut the sun and maybe even some cases in which it is clearly true.⁴⁴¹ In which case it seems that “Max cut the sun” does communicate some content even though it is a token-underdetermined content. There may be a possible world in which Max and her team at NASA are able to construct a device that splits the sun into two parts. This might count as a case in which Max cut the sun, or it might be a case in which the Max helped to cut the sun, or it might be that the device does not operate by what most would call cutting, it might have pulled the sun in two. I will discuss this understanding of content (as a partial function) in the next chapter. More importantly, there may be cases in which Max has taken a very large cutting implement, used it on the sun, and now the sun is in two pieces. That would

⁴³⁹ Ibid 225

⁴⁴⁰ Ibid 225

⁴⁴¹ Borg 2004: 236

seem to be case in which the sun has been cut. When the sun is in one piece, it is false that Max has cut the sun. In which case, it seems that there is still some content despite a lack of Background assumptions.

Searle might maintain that in each case we are not filling out a single partial function in different ways but are instead choosing different contents each of which are dependent on a Background. This, at least, is the moral he draws by considering “snow is white”. That some partial functions will have the same extension for a given possible world, as in the clear cases of Max cutting the Sun or not, does not mean that they are not distinct contents. This does not mean that the partial function view is incorrect, just that it is not automatically better than the type-underdeterminacy view. Though it may be an assumption of Searle’s that this kind of partial content is too weak to count as actual content as it may do very little to distinguish worlds into two distinct categories.

Searle would not be alone in thinking that partial functions make truth-conditions too weak.⁴⁴² On this view, ‘T-sentences [of the form ““P” is true iff P”] display knowledge of truth-conditions in that sense [of knowing which state of affairs must hold for a sentence to be true] only if the right hand side of the biconditional is *used*, that is, only if the necessary and sufficient condition which it states is transparent to the utterer of the T-sentence’.⁴⁴³ The idea would be that if one does not know when the truth-conditions are satisfied then one does not know what those truth-conditions are. So, one does not know what the representation means. This might be especially important when it comes to thoughts, as we want to know what might follow from a given thoughts being true or not. But we can only tell what follows from a thought being true if we know how things are when they are true.⁴⁴⁴

This argument can be taken to apply to cases in which we have a partial function if we have reason to believe that one might be unclear on either the necessary or the sufficient conditions. In the case of

⁴⁴² This line of response is suggested by Recanati (2004: 93),

⁴⁴³ Recanati 2004: 93

⁴⁴⁴ Harman 1999: 200-201

“Max cut the sun” we know it is a necessary condition that the sun be in two pieces and that Max be somehow causally responsible for this.

What is not clear are the sufficient conditions. Is it enough that Max worked as part of the team at NASA, or pressed the switch, or will she have had to wield a cutting implement? It seems that these are legitimate questions left open by the sentence and it is not clear how to answer them based on the sentence alone. So even a view of partial functions can be affected by this argument as the sufficiency conditions are left open. If partial functions are too weak to count as genuine truth conditions, then it seems that Searle’s argument goes through here. “Max cut the sun” may be type-underdetermined, and there may not be a way of making it determinate. If there is not a way to fix these conditions, as some contextualist positions can, then this may well seem to be too weak to count as content at all.

If partial functions are acceptable as truth-conditions, then there is an acceptable minimal content. In which case, the sentence is not type-underdetermined as Searle’s argument would suggest.

However, it seems that a great number of thoughts would be token-underdetermined on this approach. For instance, the thought that “James is tall” or “Wales is big” might be cases which would also become partial functions. The reason for this is that whilst there are going to be cases in which someone is clearly tall, ‘there seems to be no non-arbitrary way to settle on a candidate for what the invariant meaning should be’.⁴⁴⁵ In this case, it is not obvious what the necessary and sufficient conditions are for being tall or big outside of a contextually determined parameter. There may still be cases in which something is clearly big, say, the universe, and cases in which something will always fail to be big, e.g. the smallest sub-atomic particle. In between those points it is not clear whether something will count as big or not independently of some context e.g. big for a country or big compared to mammals.⁴⁴⁶ This suggests that if one is going to take a route that allows for partial

⁴⁴⁵ Jackman 2007: 334

⁴⁴⁶ MacFarlane (2007) seems to have something like this in mind when suggesting that Minimalists should instead be relativists.

functions then there should be a way of getting to a total function, or at least to determinacy in that case. Minimalists seem to be able to do this by relying on an additional content, such as implicatures making use of more substantial non-linguistic contents. However, for reasons that I will discuss below, the same option is not available to proponents of the Mixed View (MV).

For proponents of the MV, the challenge is to find a way of claiming that we can entertain a thought with a content that we are not aware of (assuming that Searle is right that one and the same thought can express different contents). The Mixed Viewians cannot make the same move that the Minimalists would. This is because the Minimalist view makes use of pragmatic factors in communication that cause us to focus on an interpretation of an utterance at the expense of attending to the minimal content. For proponents of the MV, however, there is no pragmatics of thought.⁴⁴⁷ In which case, there is no distinction between the content that the thought has in virtue of its literal content and the content that it has in virtue of any pragmatic work. On their view, the content that we are aware of is the only content that that thought has.

So, it is not clear how we would be able to miss the content that that thought has, given that the thought does not have to be interpreted in the same way that utterances are. In which case, proponents of the MV are not able to make the same moves as the Minimalist as they cannot appeal to a further pragmatic element to thoughts that we are aware of at the expense of a minimal content.

In which case, it seems that Searle's arguments push the MV into a dilemma according to which there is either no content for a given thought and thoughts are type-underdetermined or the content is only partial, and thoughts are token-underdetermined. In the following chapter I will investigate the possibility of token-underdeterminacy so in this chapter I will proceed as though thoughts are type-underdetermined. If type-underdeterminacy proves to be a possibility, then it will help to undermine the MV.

⁴⁴⁷ Fodor 2005: 104, Carston 2002: 76

The way for a proponent of the MV to resist Searle's view is to argue that there is no easy inference from the features of natural language to the features of thoughts. I will discuss the inference from natural language to thoughts below, but this approach has the potential to backfire on proponents of the MV who want to make use of some of the similarities between thoughts and natural language sentences by positing a language of thought. So, they need a motivation for the view that the two are quite different. A key motivation for this view is the regress argument. Proponents of the MV are likely to insist that there is no way that a pragmatics of thought does not lead to a regress and that the burden of proof is on RRCs to show how this can be avoided. I will now consider Perry's concerning view.

4.3 Perry's Concerning Relationship

In this section I will elaborate on the concerning relationship and show how it can be used in understanding Searle's position in a way that can avoid a regress. The claim will be that the Background can cause a thought to concern a way the world might be by narrowing the range of possible contents for that thought down to one.

Perry distinguishes two ways in which a representation can stand in a relation to the world. The representation can be *about* something or it can *concern* something. Perry describes the relations as follows: 'We shall reserve "about" for the relation between a statement and the constituents of its content, articulated and unarticulated. We shall say a belief or assertion *concerns* the objects that its truth is relative to'.⁴⁴⁸ There are a few ways in which one might interpret this distinction. The interpretation that I will use here may not be exactly the one that Perry had in mind, but this exegetical question is not my primary concern here. One way of understanding the aboutness versus concerning claim is that there are a couple of tiers of significance that a representation might have. These are character and content (see chapter one). Initially, it seems that aboutness comes in at the level of character, whilst concerning involves a relation to content. I will elaborate on each in turn.

⁴⁴⁸ Perry 1986: 147

What a representation is about is determined by its character. I propose that this is best understood as a relation between the representation's logical form and its character. The logical form of a representation is also 'a well-formed formula, a structured set of constituents, which undergoes formal logical operations determined by its structure'.⁴⁴⁹ All the parts of the logical form of a representation go into determining the character of that representation and there is no part of the character of a representation that is not determined by the logical form of that representation. This means that any unarticulated constituents that do not appear on the surface level of the representation (e.g. the words in the sentence) but are in the logical form of that representation will be reflected in the character.

A representation concerns a state of affairs in virtue of having a content. What a representation concerns can be found by considering several distinct possible worlds. If the extension for a content varies between two possible worlds when there is only one difference between those possible worlds, then we can tell that that content would put a representation in a concerning relation with the factor that was changed. For the thought *THERE'S A CAT ON THE MAT* and we consider a possible world with cats and mats but no cat on a mat the content will give the value false. If we consider a world that is the same as the first except that now at least one cat is sitting on mat, then the extension will equal true.

Importantly, the content need not be determined by the character on this view. The character of a representation might serve only to limit the range of possible contents available. So, there can be type-underdeterminacy on this view. It is also significant that the determination of content can be sensitive to factors that do not appear in the logical form of the representation and are therefore not a constituent of that representation. This means that a representation can concern things that are not a constituent of that representation. It can concern things that it is not about.

⁴⁴⁹ Sperber and Wilson 1995: 72

There may be cases in which this range of contents gives the same truth-value as an extension for a given possible world, so there may still necessary conditions on the truth of a representation before a given content is assigned to it. For instance, all true instances of “There is a cat on the mat” will require that there be cats and mats. This can account for our intuitions of necessary truth-conditions without being committed to partial functions for content.

Below is an example to spell out the difference between aboutness and concerning, but the important thing to note from the outset is that the concerning relation does not require additional elements to be added to a representation in either its surface form or logical form. This is significant as it allows proponents of type-underdeterminacy at the level of thought to avoid a regress. The regress can be avoided as additional elements do not need to be added to the representation to change what the representation concerns. Instead, the representation can remain fixed whilst the content that it is in relation to changes as the result of changes to some further factor. The current proposal is that this additional factor is a Background of assumptions.

To illustrate the distinction between what a representation concerns and what it is about, and how the concerning relation might change depending on the Background, consider the Z-landers:

Consider a small isolated group, living in a place we call Z-Land. Z-Landers do not travel to or communicate with residents of other places, and they have no name for Z-land.

When a Z-Lander sees rain, he will say to others not in a position to look outdoors, “It is raining”. His listeners then act appropriately to there being rain in Z-land: they close the windows in Z-land, cancel plans for Z-land picnics, and grab umbrellas before going into the Z-land out-of-doors. They have no other use for “It is raining”. They do not call their sons in far-off places, or listen to the weather news, or read newspapers with national weather reports.⁴⁵⁰

⁴⁵⁰ Perry 1986: 144

Here the idea is that the Z-landers' utterances and beliefs are not about Z-land, when they say or think that it is raining. Z-land does not enter the thought as an unarticulated constituent or in any other way. It is not present as a part of the character of that utterance. Instead, their beliefs concern Z-land. As Perry puts it 'There is no need to postulate a concept or idea of Z-land as a component of their thought in order to secure the connection to Z-land. The connection is secured by the role of the whole belief in their lives'.⁴⁵¹ This is important as it suggests that a belief can concern things that it is not about.

This also suggests how the concerning relation might play a role in explaining how the Background determines the content of an utterance. The Background includes a set of assumptions that make it possible for these representations to play a role in a person's life. A worry that one might have with this is that to play a role in a person's life, by letting them understand the content and how to act on it, that character alone is insufficient for this interaction to take place. It is not determinate enough. In which case, the representation will already have to have a content assigned to it. So, the representation was not type-underdetermined to begin with and (b) is false, where (b) is the claim that all (or almost all) thoughts are context-sensitive.

In more detail, the worry goes as follows. The thought comes with a character, but not a determinate content (by hypothesis). To determine content, that thought needs to interact with Background assumptions. There is a question as to whether the character could be sufficient to have an interaction that can play a role in someone's life. This is because character is, on this view, quite a thin notion that only serves to limit the range of possible contents. If the character is insufficient then an alternative is that it is the content of that representation's parts that allows it to interact with the Background and helps an agent to know what to do with that information and to assess the truth-value of that content. In which case, the thought already had a determinate content prior to its interaction with the Background.

⁴⁵¹ Ibid 144

This worry can be avoided. To show this, more detail is required to explain how character can interact with Background assumptions to select a content. Here I intend to use Relevance Theory as a way of explaining this. One way that this might work is that the content that is most relevant, given the manifest Background assumptions and the constraints imposed by character, is the content that is selected for that thought at that time. A content is relevant to the extent that it produces cognitive effects (usually in the form of affecting assumptions that one has) and irrelevant to the extent that it takes cognitive effort. The content that is most relevant is the one which has the most effects for the least effort.⁴⁵² This can change depending on which assumptions are manifest. For instance, if it is highly manifest that I am in Z-land then the content that *it is raining in Z-Land* will have the greatest effect for the thought IT IS RAINING. If it was manifest to me that I am in Portugal, then that *It is raining in Portugal* would likely be the most relevant possible content. This does not require that the thought already had a content.⁴⁵³

The advantage of understanding Searle's view as that of a Background of assumptions causing a given representation to be in a concerning relation with a state of affairs by selecting a content is that it does not require an additional thought to be generated. Instead, one of a range of possible contents is selected based on which of them is the most relevant to an agent based on what is most manifest to them at that time from their Background. The thought itself, understood as a mentalese sentence, does not need to have additional concepts added to it. What it is *about* need not change. In which case, a thought can have its type-underdeterminacy removed without requiring there to be an additional thought that attempts to make it determinate. In which case, RRC can avoid a regress

⁴⁵² See Carston (2002: 45), Sperber and Wilson (1986/1995: 125), Relevance theorists take this process to apply to ostensive-inferential stimuli that are used in communication. Carston (2002: 76) also considers whether the Background might play a similar role at the level of thought but seems to reject this as a possibility, largely on the basis that it would trigger a regress. If, however, the content can be selected without requiring an additional thought, as the concerning relation suggests, then this is a possible view.

⁴⁵³ Many acts of communication can gain content in virtue of the context in which they are performed. For instance, Relevance Theorists Sperber and Wilson (1986/1995: 48) argue that the act of leaning back on a bench can be an action that communicates a proposition to another person e.g. that there is an ice cream vendor or that a dreadful bore is on their way over. So, this act communicates a content but the action type of leaning back on a chair does not have a content.

from occurring at the level of thought. If the Background changed what a thought was about then there would be a regress if one accepted RRC. In the next section I will consider some objections to this proposal.

Before turning to a discussion of some objections I will contrast the view that I want to defend here with semantic relativism. On a relativist view sentences in natural language will express a constant content but that content will only be truth-evaluable relative to contextual parameters beyond the possible world under consideration.⁴⁵⁴ For example, there may be a requirement of relativizing a content to an individual's taste for sentences of the form "x is tasty".⁴⁵⁵ On the view I am proposing here the character expressed is constant but the content can vary depending on the Background. So, it is the character that requires additional parameters to determine a content not the content itself.⁴⁵⁶

4.4 Objections

We have seen above that a RRC may want to claim a given representation can only *concern* a proposition with the help of a Background and that this move avoids a regress. That does not mean that the resulting interpretation of Searle's view is free from objections, however. Some of these I will consider below. The objections that are most troubling are those which suggest that this view is impossible or ones according to which the cost of accepting this solution is too high to be palatable.

4.4.1 A Slippery Slope?

To begin with, whilst this approach would keep the Background theorist from having to continually add more to a given representation and thereby preventing a regress, it seems that they may still be subject to a sort of slippery slope argument. The difference being that there is no clear stopping point for the number of parameters that one needs to determine a content. This is a problem as the

⁴⁵⁴ Kölbel 2009: 375–376

⁴⁵⁵ Ibid

⁴⁵⁶ There may also be some overlap between Searle's Background theory and what Cappelen (2008:23) has called content relativism, where 'An interpretation sensitive sentence can have one content relative to one interpreter and another content relative to another interpreter'. How exactly these views relate to one another is not important for the rest of the thesis.

Background is vast and all of it can be manifest to an agent to varying degrees at a given time. If the entire Background is affecting each thought, then the agent must do an implausible amount of processing to derive a content for a given thought. So, there is some pressure to delineate which Background assumptions are the ones that have an effect on a given occasion.

How to decide which Background assumptions to use on a given occasion and the influence that they should have when determining content is not obvious. A natural first try would be to decide along the lines of what is most manifest at that time. However, that suggestion runs into problems when considering trivial cases that are easy to represent. Sentences of the form "If P then P" are both very likely to be accepted as true and easy for an agent to represent, making them very manifest at a given time. We don't want agents to have to be considering representations of this form every time they entertain a thought of another form, yet it is not clear how these cases can be ruled out based on manifestness. What's worse, is that it is not clear how these kinds of trivial assumptions would help to determine the content of a representation. It seems that what is usually required are non-trivial assumptions about what it is to be a gardener or a cake cutter or one's location. There is no guarantee that these non-trivial assumptions will be the ones that are most manifest as they may be more complex to represent for an agent than trivial truths (this does not mean that they are difficult for an agent to represent, just that trivial truths may well be easier and therefore more manifest). Non-trivial assumptions may require more concepts to be used in a more complex configuration as well as being harder to accept as true.

Perhaps one could say that this is a misunderstanding of what it is to be capable of representing an assumption. Perhaps it is not a case of how complex the representation is but more a case of how likely it is that the agent would represent that assumption at that time. This stretches the word "capable" but some passages on manifestness seem to suggest that something like this is the correct understanding. For instance, the degree of manifestness of an assumption is described as something that 'may shift from moment to moment depending on features of the external physical environment

and on [the agent's] internal cognitive states (for instance, where attention is focused)' and that other assumptions will be manifest to someone only for a short time depending on '...a function of where one happens to be who one happens to be with, etc.'⁴⁵⁷ This suggests that the perceptions one has of one's immediate environment play an important role in determining what is manifest to an individual at a given time. So, a subset of what is manifest to an individual will be determined by what that individual perceives in that environment. Further assumptions may then be made manifest to an individual based on their relevance (in the technical sense suggested by Relevance Theory) to the situation that the agent perceives themselves to be in. On this picture, what is manifest to an individual is either what is perceived by them or what is relevant to them. This provides a way of ruling out trivial assumptions from influencing the content of a given representation when trivial assumptions should not have such an effect.

There is a worry that there is some circularity here. Manifestness is defined in terms of what one could accept as probably true and be capable of representing to oneself. What one is capable of representing is determined in part by relevance. Relevance, however, is calculated (at least in part) based on what is in one's Background. If so, what is manifest is determined by what is manifest. That is a problematic loop to be in. In response, some assumptions in the Background will be determined based on what is relevant at a given time. Some will be determined based on what one perceives. This subset based on perception provides a basis for determining the relevance of other assumptions in a way that does not require making use of things that were already relevant.

In which case, all assumptions that are manifest to an individual at a given time are manifest either because of what one perceives or because they are relevant. At least some of the assumptions that are relevant will be so because of their relation to what is perceptible. From that point assumptions can continue to be generated based on previous assumptions until a point of irrelevance is achieved (if there are no interruptions from the environment).

⁴⁵⁷ Carston 2002: 68

For example, one might perceive a cat, and the assumption that there is a cat becomes a part of the Background. This may cause typical cat behaviours to be relevant, adding to the Background. This may trigger further assumptions to become relevant, such as one's previous encounters with cats, that they make good (or bad) pets, that they can damage ecosystems, that ecosystems should be protected and so on. These further assumptions become decreasingly relevant as more work is required to entertain them or as one's perceptible environment changes. Whilst this does not altogether remove the possibility of the Background determining itself, it does so in a benign way. If so, then useless but trivial assumptions would not be the most manifest as these would not be useful in the pursuit of determining content.

The problem with this approach is that it moves the question rather than answers it. How is it determined which assumptions are relevant in a given scenario, so that they can be suitably manifest? Even tautologies might be useful on occasion. In the above example, a range of other assumptions could have been made manifest, for instance that there are big cats or that dogs make good pets as well. It is possible that all these assumptions become manifest to some degree. How to determine which assumptions become most relevant is not obvious and may depend on other factors such as a person's preferences, emotional state and other idiosyncrasies. Manifestness was being used to determine which assumptions would be the ones that are used but it now seems that manifestness cannot answer the question, instead relying on some other notion of relevance and probability. Even on a well-known understanding of relevance, as found in Relevance Theory, it is not immediately clear how a cognitive process would determine which assumptions meet these criteria.

This implementation problem serves to make the concerning relation between a representation and its content seem mysterious. It is not clear how the Background is meant to interact with the representation because it is not clear how much of the Background and which bits of the Background are involved in this interaction. Stating that it is the relevant assumptions that play that role is not a huge advance in this regard. So, a proponent of the MV might want to claim that although this

understanding of RRC might still be consistent, it comes at a high price. Namely, that it is now not at all clear how this view should work. In which case, it is not really a solution and should be disregarded.

To respond to this, I would like to raise two points. One is that the concerning relation is not as mysterious as it seems, despite it being unclear how much of the Background is in play at a given time. The other is that the proponents of the MV do not fair that well in this regard either. It is a high price to pay but it is one that we are already paying due to the frame problem (see chapter 2). I will take each in turn.

Whilst it is not clear how much of the Background should be used on a given occasion or which Background assumptions should be used, this does not make the concerning relation itself so unclear. The concerning relation is one on which a representation is assigned a given content. This occurs as the (let's suppose) non-indexical character of that representation can give a content as its output. What is interesting about this relation is that the output is not only dependent on the character and the context of use, but on the Background of the individual who is tokening that thought. The Background acts as an additional parameter here. Whilst there are questions about how the value of this parameter is determined, that does not seem to make the concerning relation itself so unclear. How much of the Background is used will likely be an empirical question.

Another point to make here is that whilst understanding how some Background assumptions are selected on one occasion and not others is difficult, this difficulty does not seem to be so different to one of deciding which premises to make use of for a given inductive or abductive inference. Both problems depend on there being some form relevance and both views struggle to say how a cognitive process can distinguish relevant from non-relevant information. It is also worth noting that the concerning relation is not likely to be carried out by a modular process, as the Background is domain general. The point being that whilst this obscurity is a high price to pay, it is a price that we are already having to pay because of the Frame Problem.

4.4.2 A Flawed Analogy?

Another objection that a proponent of the MV might want to make is that I have not proven that a type-underdeterminacy view is correct. That is true, but it was not my aim to prove that RRC was correct. My aim was to show that it is possible to maintain RRC without leading to a regress.

Understanding Searle's view as involving a concerning relation is a way of showing a possible RRC view without the risk of a regress. Thus, I would suggest that my aim has been met. There are, however, some other worries in this area. One is that it is not possible to have context-sensitive thoughts not because this leads to a regress but because thoughts are the basis for all other forms of intentionality, such as that displayed by natural language sentences. On the view that I am proposing here this may not be a possibility. A thought by itself no longer has intentionality. It does not have a content in the absence of a Background. So, the thought no longer seems to be the basis of intentionality, as that would seem to require a content.

Whilst the regress argument seemed to be a significant reason for thinking that thoughts were importantly distinct from natural language sentences, it may be that there are other motivations for wanting to keep thoughts as the basis of intentionality. For instance, thoughts being the basis of the intentionality of natural language sentences might be a useful way of distinguishing those symbols that only arbitrarily have a certain content from those that do not. The word "dog", they might say, has the content that it does because of its relationship to the concept DOG. Nothing about the word "dog" forces this on us, "hund" may also have been used and "dog" may have meant *flower*. However, the concept DOG is not arbitrarily connected to *dog*. The Background view proposed here seems to put this distinction in danger as the content of DOG now seems to depend on the Background when it is tokened to have a given content. That, they may object, is too much arbitrariness for thoughts or concepts to have.

However, I would argue that there are still ways in which this distinction can be respected on the Background view that I propose here. It may be that whilst the connection between "dog" and DOG is arbitrary, the connection between DOG and its possible contents is not. In this case, it is the

character of the concept DOG that is a non-arbitrary part of that concept. That character may limit the range of possible contents for that concept so that it is always about (in Perry's sense of the term) dogs. However, for it to concern dogs requires a Background.

A proponent of the MV might also argue that there is no way to prove that the Background view is correct, or even to motivate it. The objection might be that there is no straightforward inference from natural language sentences' context-sensitivity to the context-sensitivity of thoughts. The reasoning may be that to do so would be to assume that the two kinds of representation are analogous in these respects. However, any such attempt to do so would depend on assuming what is to be proved. In which case, any attempt to argue in this way would be question begging. Especially as proponents of the MV seem committed to denying that there is any such analogy in virtue of maintaining claim (a), that all (or nearly all) natural language sentences are context-sensitive, whilst denying (b), that all (or nearly all) thoughts are context-sensitive. A crucial difference, that suggests a disanalogy, is that natural language sentences and their context-sensitivity, are due to their use as tools for communicating with other people, whilst thoughts are not.⁴⁵⁸

There is some truth to this. However, assuming an analogy between natural language and thought is not so unreasonable as it might appear. Both are systematic and productive, for instance. Both can also be assigned truth-conditions and neither depends on factors like resemblance to do so. This makes them more alike than other forms of representation, such as painting which does (at least often) depend on resemblance to represent things and is not systematic or productive. So, it does not seem unreasonable to suppose that thoughts and natural language sentences might also be similar in other respects.

Furthermore, whilst the pragmatic work that must be done on natural language sentences when they are uttered might be due to their use in communication, this is not the only explanation of their context-sensitivity. It may be that context-sensitivity is just a feature of the sorts of representations

⁴⁵⁸ This seems to be the view of Carston (2002: 76)

that humans tend to employ. It may be that the representations that humans tend to employ need a degree of flexibility to be useful. Being able to select a relevant content for a concept would be a good way of taking this into account. It means that the content accessed is likely to be a useful one where possible.

A similar role is played by *ad hoc* concepts on the MV account. An *ad hoc* concept is one that is composed for a specific purpose, usually a variant on an already existing concept but with an important difference to its content.⁴⁵⁹ The Background view does not require the composition of additional *ad hoc* concepts in order to account for flexibility in human representation. This serves to make the Background view more parsimonious than an *ad hoc* account.

This does not vindicate an argument by analogy, but it does show that it is not wholly unreasonable. If it is correct that the regress argument is not a reason to think that the two must be dissimilar, then it becomes more plausible that they are similar. Furthermore, where natural language has linguists to study it, the prospects of a similar field existing for thoughts seems less likely, if only because thoughts are much harder to access than natural language sentences.⁴⁶⁰ In which case, dismissing the idea that there could be analogies between the two makes it unclear how to understand thoughts. It is correct that any such induction will be weak, but it is not completely out of the question.

4.4.3 Holism

A third concern with the Background view is that it commits one to holism about content and there are several important objections to holism. This view begins to look like holism as the content of a given thought is dependent not only on the thought itself but on what is manifest to an agent. The worry with this sort of a view is that the content of a concept will change whenever what is manifest to a person changes i.e. whenever their Background changes.⁴⁶¹ This makes agreement between

⁴⁵⁹ See Allott and Textor (2012) and Wilson and Carston (2007) for more on *ad hoc* concepts.

⁴⁶⁰ During the linguistic turn it was believed that thought was linguistic, so that one could infer the nature of thoughts from the nature of language. Indeed, studying language was considered to be the only way of studying thought. See, for instance, Chomsky (2006: 162) and Lievers (2005: 184-5).

⁴⁶¹ For a version of this criticism that makes use of beliefs, see Fodor and Pylyshyn (2015: 55).

individuals, and even the same individual over time, difficult. Given how much the Background can change at a given time, the content of a concept will also change (if the manifest assumptions that change or are introduced are relevant ones). To illustrate this, suppose that yesterday I tokened the thought that BUTTERFLIES ARE ELEGANT. Between now and then I have done research on both butterflies and elegance, so that the next time I tokened the thought BUTTERFLIES ARE ELEGANT it now has a different content. In which case, it appears that I no longer agree with my past self because I have a different content for each thought. A similar worry appears if I change my mind because of the research and think BUTTERFLIES ARE NOT ELEGANT. If the content of BUTTERFLIES ARE ELEGANT has changed, then so has the negation of that thought. In which case I am not able to contradict my first thought that they are elegant. The worry is not just that this is implausible but that we risk a version of relativism that makes disagreement about the world impossible. In which case, *prima facie* incompatible beliefs are in fact compatible with each other. So, we no longer contradict our previously held beliefs and they don't need to be classified as false, even if new information would typically lead us to say that those beliefs were in fact false.

One point that can be made against this objection is that it places the bar for agreement and disagreement very high. The assumption seems to be that if there are different contents then it is not possible to agree or to disagree with one's past self. Whilst there will be cases where a difference in contents will make agreement impossible, for instance, thinking at t_1 that *cats are blue* and at t_2 that *dogs are not blue* there will not be any disagreement or agreement. The two contents do not relate to each other in this way. However, it is not obvious that agreement or disagreement requires that the two contents are identical.⁴⁶² To begin to motivate this view I will first consider inter-personal cases of (dis)agreement. I will then consider intra-personal cases. If (dis)agreement can be understood as only requiring that the two contents are similar enough, then the possibility for

⁴⁶² A similar point is made by Rovane (2013: 399), though she emphasizes that people cannot help but share a great number of Background beliefs in virtue of living in the same world.

change in content and the holism that it implies is not so problematic as (dis)agreement is still possible.

One point to note is that communication between individuals does not require that they each have the same content in mind to be successfully communicating. It is enough that they each have similar contents. For instance, on Carston's view, the pragmatic processes that one undergoes when comprehending an utterance is '...aimed at constructing the propositional form intended by the speaker, or one similar enough to have the intended effects'.⁴⁶³ Here I want to focus on the idea that successful communication often only requires that the hearer entertains a content that is similar enough to the one that the speaker intended to convey. This also seems to apply to cases in which people agree with each other or disagree with each other. In which case, the standard for agreement and disagreement is lower than this objection suggests it is.

It seems that one needs to make use of a good enough view of communication when one has an ostensive inferential view of communication. On this view of communication, a speaker communicates a proposition by providing 'direct evidence of one's intention to convey it'.⁴⁶⁴ There are two intentions that a speaker has when communicating in this way. The first is an informative intention '...to make manifest or more manifest to the audience a set of assumptions I'.⁴⁶⁵ The second intention is the communicative intention to 'make it mutually manifest to audience and communicator that the communicator has this informative intention'.⁴⁶⁶ For this discussion I will focus on the informative intention as the most important for determining whether communication is successful. It seems that in those cases in which the set of assumptions I is large then communication may be successful when the hearer entertains most of them. In which case, identical contents are not a necessary requirement for communication. This, however, does not do enough as there will be

⁴⁶³ Carston 2002: 59

⁴⁶⁴ Sperber and Wilson 1986/1995: 23

⁴⁶⁵ Ibid 58

⁴⁶⁶ Ibid 61

cases in which the informative intention is more specific or more demanding than this. It may be that there is only one assumption P that they want to communicate.

I will argue that successful communication does not occur only when speaker and hearer entertain the same content. Instead, it occurs when the contents they entertain are sufficiently similar for their purposes. The upshot of this is that people can agree or disagree with each other even when they entertain different contents. If they could not agree or disagree with each other on this view, then it seems that agreement and disagreement become too difficult. Furthermore, when successful communication of a proposition(s) occurs, it seems that this makes it possible for the interlocuters to agree or disagree with each other on whether the propositions communicated are true or not. If agreement and disagreement can occur inter-personally even when two individuals are entertaining different contents, then it seems that the same can be true of the intra-personal case.

Given that natural language sentences, even when you think that they determine truth conditions with little pragmatic work, leave a lot open to the hearer of an utterance in terms of what they take to be the communicated content, it is unlikely that interlocuters often entertain the same content(s) as the one that meets the communicative intention of the speaker. (If you are a Minimalist then you may think that they are likely to entertain the same minimal content at some point during their comprehension of the utterance, but this is also unlikely to be one that is communicatively rich enough to meet communicative intentions). Given that people can and do agree and disagree with each other, despite seeming to entertain slightly different contents, it seems that a good enough approach to content is sufficient for agreement and disagreement.

For example, A may say to B “Butterflies are elegant”. They may each have subtly different contents that they associate with “elegance”. There may be some cases in which A's content would give a different output for some inputs to B's content, so that they differ in extension as well as intension.⁴⁶⁷ Nevertheless, the contents give the same output for most cases. Suppose that they both

⁴⁶⁷ Church 1941: 3

assert that butterflies are elegant. In which case, it seems that one can say that they agree with each other. For most of their purposes they can co-ordinate with each other on things that are both elegant and butterflies. The same can be said of disagreement.

Considering the case of a thought JOHN CUT THE GRASS tokened by two individuals, one of whom has a Background focused on cake cutting, the other on using a lawn mower. There would be circumstances in which the two can count as agreeing with one another, despite their having differing contents. It may be that all that matters was that some sort of cutting was performed. Here both can agree. If it matters that the grass was cut efficiently or to look well-kept then the two will disagree.

In this case, A and B may have different contents for their concept of ELEGANCE, but their behaviour will still be similar when they agree and importantly dissimilar when they disagree. They cannot have too widely divergent contents as their concepts will still have a character that constrains the range of possible contents. For instance, when they agree both will appreciate the aesthetic qualities of butterflies and may be seen to admire them. When they disagree B may not appreciate any aesthetic qualities of butterflies or will appreciate different qualities. (Supposing that B thinks BUTTERFLIES ARE NOT ELEGANT but thinks that BUTTERFLIES HAVE CAPTIVATING COLOURS then B may prefer to see static images of butterflies rather than seeing them in motion).

This does not mean that there will not be cases in which a person intends that a content P be entertained by their audience and that P is the only content that will do. In these cases, there will be a high bar on what is required for communication, along the lines that Fodor and Pylyshyn suggest. I want to suggest that this is often not the case. This is partly because forming such a specific intention requires more effort than a condition that is easier to satisfy. It also requires that the speaker does more work to make sure that they communicate that their intention is this specific. So where set I has only P as a member, it may be either that the recipient must access that exact content, in which case communication is less likely to succeed, or it is satisfactory that the hearer entertains some

similar content P^* . The similarity may be borne out by the behaviour of the hearer (to behave as A would expect them to if they believed P was communicated). The point is that informative intentions will not always be so demanding as to require that a specific content be entertained by the hearer. Often, any of a range will be satisfactory. This range could be determined either by their compatibility with the character of an intended thought or with the behaviour that they would then engage in. This latter point comes about as we often use speech as a means of co-ordinating behaviour.⁴⁶⁸

A second point to make here is that in those cases where a speaker intends that an exact content is communicated, they often must introduce the idea to their audience by giving their audience enough information about the content for them to entertain it. They may have to present a context for the content and to say what the content is not and so on. This can be understood on this model as making certain other assumptions manifest to an audience so that they appear in that audience's Background. That way, when the audience tokens their thought, they are more likely to be entertaining the exact content.

It is worth emphasising that holism can make use of Background assumptions to help holists to determine the content of an utterance and this can allow for disagreement as well. Consider someone who sincerely utters "Trees grow on Mars". We can still attribute a false belief to this person if we press them further and find that they also believe that trees are a kind of plant, that Mars is a planet other than Earth and that growing is a process by which organisms become bigger. At which point, it seems that we must attribute a false belief to the speaker.⁴⁶⁹ In which case, disagreement is possible even on a holistic approach. This is so even when the speaker intended to communicate something specific.

⁴⁶⁸ See Guerts (2019) for a development of this approach to speech acts.

⁴⁶⁹ Rovane 2013: 398

An opponent might object that there is no principled point at which to stop this investigation into the speaker's Background assumptions. Won't we also need to know that they take planets to be a celestial body in orbit around a star and what they take stars to be and celestial bodies and so on.

There might always be a way of avoiding having to attribute false beliefs. In response, it is important to bear in mind that Background assumptions need not themselves be representations or represented. They are not themselves thoughts, at least not in the sense that I am interested in. This means that it is possible that their connection to things in the world operates in a different way. It is open whether they depend on further Background assumptions or not. If not, then there is not a problem here. If so, then there are relevance problems again. These can be responded to as above.

Furthermore, two individuals A and B may have different backgrounds that affect the content that they have for their thought that IT IS RAINING (I am assuming that they can have the same concepts as the concepts can have common characters). Nevertheless, they can still agree and disagree about the weather. Suppose that A is an inhabitant of Z-land and B is in London (and not from Z-land). Here it seems that even though the contents that they will have are different, and this difference is due to their Backgrounds, they can still agree or disagree with each other. B can agree that it is raining in Z-land, though she may need a different thought to represent that, or a change to her Background to do so. Nevertheless, she can access the same content that A can. Similarly, disagreement can still occur here, though they may have to token different thoughts for this to come about. So, where A thinks IT IS RAINING, B would have to think IT IS NOT RAINING IN Z-LAND. Whilst this is not as elegant a picture of disagreement, as it is possible for A and B to token the same thought and still disagree, this is something that we should expect of a view on which the content of a thought can change.

When it comes to the intra-personal case, similar moves can be made. One is to claim that for an individual to agree or disagree with themselves across time does not require that they always be entertaining the same content. As communication across individuals does not require interlocuters to entertain identical contents and these individuals may still agree or disagree with each other, it

does not seem that agreement requires identical contents be entertained by the speaker and the hearer. It is not clear why this should be different in the intra-personal case. Another is to claim that for that (dis)agreement to come about requires the individual to think different thoughts to get matching content. Either option seems to make the worry with a holistic approach less worrying.

A related concern with holism is that it leads to relativism, so that one's beliefs do not need to be changed considering new evidence. The idea would be that the belief (or the belief's content, in this case) remains untouched as the evidence, or contrary content, comes from a different Background. In which case, the content of belief is determined in such a way that other beliefs cannot touch it. Communication aside, this seems to be another motive for finding holism problematic.

There is, however, an argument against relativism on holistic grounds. The target of this argument is a version of relativism according to which there can be alternative Backgrounds that are true but not translatable into one another.⁴⁷⁰ As they are not translatable into one another, there can be no agreement or disagreement. This is relevant to the intra-personal case if a person's Background could change such that their Background at t1 makes contents cease to be translatable into the person's beliefs at t2 on the basis that their Background has changed so much between t1 and t2. In which case a person may have two thoughts that are true relative to their Backgrounds at that time, but which cannot affect one another's truth-values. This means that the truth or falsity of one cannot affect the truth or falsity of the other. In which case they cannot be incompatible with one another. So, agreement and disagreement cease to be a possibility. The argument against relativism aims to establish that we cannot help but share certain Background assumptions which make it the case that either the beliefs are translatable, or they are nonsense.

The argument proceeds on the basis that we all occupy one world, where the world is understood as all that is the case.⁴⁷¹ If we occupy the same world then we are all likely to share an important set of

⁴⁷⁰ Ibid 397. Rovane uses "conceptual scheme" instead of "Background" but this should not affect the argument too much.

⁴⁷¹ Ibid 402

Background assumptions. These assumptions form a core of assumptions on which one can build other more controversial assumptions. If it is possible that we agree on these core assumptions, then other assumptions can be in a logical relation to our own i.e. they can influence the contents of those occurrent thoughts.⁴⁷² This includes cases in which someone exists at both t1 and t2. This logical relation takes the form of compatibility or not with the set of Background assumptions. This relation makes it the case that we can understand the person in virtue of agreeing on at least certain core assumptions.⁴⁷³

The idea is that alternative Backgrounds, which would not allow for translatable contents, are not possible for humans to acquire. There is a guarantee in having inhabited the same world that any one person (or two) will have enough Background assumptions in common that their thoughts' contents can stand in a logical relation to one another. Even if one disagrees with this, it suggests that any alternative contents would be nonsense to us. We could not understand them or relate them to our own.

There are a range of cases one might encounter here. For one, suppose that someone has been raised by a cult and has come to believe that dogs are robots placed in homes by an evil demon. In this case we can understand the other person's belief. We will likely just believe that it is false, dogs do not have the property of being robots. We may also question what an evil demon would be and deny that there is a demon dedicated to this task, or that there is any way for the demon to implement its strategy effectively. We may even be astounded that there is a person who could genuinely believe that this is the case. For each of these beliefs, however, we can put them in a logical relation to our own and decide that if ours are true then theirs must be false.

A belief that is closer to nonsense may be COLOURLESS GREEN IDEAS SLEEP FURIOUSLY. This is a belief that is difficult to make sense of, at least when taken literally. We may think that it is false that

⁴⁷² Ibid 404

⁴⁷³ Ibid 404

ideas sleep at all or can be coloured. This suggests that there is some logical relation to our own beliefs still. However, it may be that we have no idea what would be the case if this belief were true. In which case it is closer to nonsense.

For a belief to be genuine nonsense on the view proposed here it needs to be that we cannot tell what the world would be like if the belief is true or false. In this case, we cannot place the belief in relation to our own. Otherwise the belief can usually be regarded as false or probably false and still put into relation to our own beliefs.

Reducing alternatives to nonsense might just be what the opponent of holism is worried about. If we cannot even understand these contents, then they cannot challenge our beliefs. Therefore, our beliefs are still protected from potentially contrary beliefs. Therefore, beliefs are unreasonably insulated and we cannot agree or disagree with someone who does not share our Background.

In response to this it is worth emphasising that the Background assumptions that form this core are difficult to avoid having. They are assumptions that one should have to successfully navigate the world. So, individuals should have certain beliefs in common, as should an individual across time. These would likely be beliefs along the lines of, for instance, concrete objects have weight, travelling takes time, there are day/night cycles, that adding more objects to a set makes the set larger and similar beliefs. These kinds of beliefs are not going to be hugely detailed or show a great deal of understanding of science, for instance. They are more likely to compose a sort of folk version of various bodies of knowledge. If one fails to have these kinds of beliefs then it seems that they have failed to engage with the world sincerely or will have engaged with it in a very different way (they may also be very young, I am ignoring this case for the moment). A different way of engaging with the world may be exemplified by someone who is kept in a padded cell from birth and kept on a constant supply of a hallucinogenic drug. In this case they may fail to acquire many of these beliefs and instead acquire some radically different ones. In which case, the prediction that what they have thought is nonsense seems correct. It also becomes mysterious how a person's Background could

change so drastically over time. So, in the intra-personal case one should not be able to change so radically. This “deep” Background ‘...is too extensive and too detailed and too interconnected to be substantially revised, and this is so even though it may take on a different significance when viewed from the perspective of one or another theoretical perch’. So, relativism is not an automatic consequence of endorsing holism.

4.5 Conclusion

In this chapter I have considered the possibility of type-underdeterminacy at the level of thought. To do so I have considered a view according to which thoughts do not determine a content without a Background of assumptions. To make sense of how Background assumptions can do this I have considered the concerning relation as a useful means of explaining how this occurs. The Background does not cause a thought to have more concepts added to it. Instead, it functions more like a set of parameters that go into determining the content of a thought, alongside its character.

This avoids the regress from Chapter Two in the following way. It avoids the need to consider an additional thought as the foundation of the initial thought. By avoiding the need to generate additional thoughts whilst taking into consideration the context as a determining factor of that thought’s content, a regress is avoided. Recall that the regress required an additional thought to be generated in each case. That no additional thought needs to be generated means that there is no regress. This makes this understanding of the Background view a useful candidate for showing that it is possible for foundational thoughts to be context-sensitive.

I then considered some objections to this view. The first was that it is not clear how the relevant parts of the Background could be selected in a given case. Relevance, or expected relevance, is a difficult property to determine and it is unclear how this property should be assigned to different Background assumptions. Whilst this is an important point, it is not one that puts the RRC in a worse position than the proponent of the MV. If the frame problem is, as I argued it was in chapter 2, a problem, then determining relevance is not something that proponents of the MV are able to determine either.

The second objection was that I have not shown that this view is correct, there is nothing to stop a proponent of the MV from continuing to be a proponent of the MV. That is correct, but I have removed one of their key defences against RRC which was that RRC leads to a regress. In which case, it is not correct that the MV has a monopoly on understanding thought.

Another worry was that it would not be possible to prove that thoughts are context-sensitive, at least not based on natural language sentences. The worry would be that attempts to do so assume what they are trying to prove, namely, that natural language sentences and thoughts are analogous in that respect. Whilst that is a problem, it seems that a weak inductive case can be made for the two to be similar, given that they share several other features.

The third objection was that this understanding of the RRC view leads to holism which makes agreement and disagreement, even for the same person at different times, impossible. In response to this I argued that this objection both requires too stringent a view of what it is to share content and that there are ways in which an agent can still access the same content across time. It is just that this will sometimes require an agent to token a different thought. This does not seem like it should cause a problem for a view on which a thought does not determine its content without some further factor. It also does not seem that holism entails relativism in a problematic way automatically.

This suggests that a view according to which thoughts are type-underdetermined is possible, contrary to the regress argument. As these thoughts do not require an additional thought to make this possible, as Background assumptions do not need to be made manifest or ever become representations to have an effect, these thoughts are foundational. In which case it is possible that there are foundational thoughts that are context-sensitive. Given that the Background view encompasses (almost) all thoughts, it seems that it is possible for context-sensitivity to be a general feature of thoughts as well. Whilst this view still has details that need to be worked out, such as how certain Background assumptions are the ones that cause that thought to concern that content, there is a comparable problem on the part of the MV as well. In the following chapter I will turn to consider

whether a token-underdetermined view is also a possibility when considering whether context-sensitivity is possible at the level of thought.

5. Travis and Judgement

In this chapter I aim to present and defend a Travis-like view of representation, according to which representations in natural language and thought are context-sensitive. That is, a view on which both (a), that all (or almost all) natural language sentences are context-sensitive, and (b), that all (or almost all) thoughts are context-sensitive, are true. These are the claims that all, or almost all, natural language sentences and thoughts are underdetermined. However, the view presented here will not be Travis's actual view. In fact, there will be some points at which it diverges from Travis's views. For instance, Travis makes no attempt to show how his view might be compatible with a computational theory of mind (CTM) (often doing the opposite). The point of this chapter is thus not to present an exegesis of Travis. Instead, it is to take insights from Travis's work and present them as a way of responding to the Mixed View (MV). I will be reading Travis as though his focus is on token-underdeterminacy, which is the claim that we only get a partial function from possible worlds to truth values from a thought. For ease of exposition I will still refer to this as the Travis view, and there are certainly strong similarities. Travis is particularly extreme in the extent to which he takes thoughts to be underdetermined in this way. To make this presentation of Travis's view clearer I will be using Fodor's position as a foil. Fodor is useful as a foil because he accepts that natural language is underdetermined whilst vehemently denying that thoughts are underdetermined.⁴⁷⁴

I will not, in this chapter, try to prove that Travis's view is the correct one. The aim is rather to present a version of his view and show that it is at least feasible as it can avoid the regress argument from chapter two. I will be presenting some of Travis's arguments for his position and I will try to defend some of the claims that Travis makes as well.

The structure of this chapter is as follows. In section 5.1 I will represent Fodor's view. In Section 5.2 I will present a rule-following argument that Travis gives to try and motivate a move away from Fodor's position. In section 5.3 I will present Travis's judgement dependent view, including a

⁴⁷⁴ Fodor 2001

discussion of some of the prerequisites for successfully judging whether an object counts as an instance of a concept. In section 5.4 I will consider how this view can respond to the regress argument. The main point of this will be that judgements do not result in new thoughts (or they do not result in thoughts that are underdetermined in a way that requires resolution at that time). The regress argument requires additional thoughts to be generated to account for context-sensitivity. I will argue that the Travis view avoids a regress because it avoids the need to resolve underdeterminacy by creating an additional underdetermined thought, which would then require an additional thought and so on.

5.1 Fodor

Fodor's view of thought has already been touched upon in chapter 2. Here I will briefly recap some of these details before including additional detail to bring out the contrast with Travis. This concerns how the extension of a given concept is to be determined. The main point of contrast between these two is in what they take the primitive bearers of intentional content to be: for Fodor it is mentalese sentences, for Travis it is acts of judgement. This does not mean, however, that the two have nothing in common, as I will suggest that they both make use of innate capacities in determining an extension.

5.1.1 Recap

I will begin by recapping some of the important details from chapter 2. Fodor accepts a referentialist view of concepts, so all that there is to the content of a concept is its referent. In Fodor's terms, 'The content of a concept is its extension; the content of CAT is the things belonging to the set of (actual or possible) cats'.⁴⁷⁵ There is no additional meaning or intension, although the "vehicle" for the concept does some of the work that intensions normally would. More generally, 'tokens of beliefs, desires, and the like are tokens of relations between minds and mental representations; that mental representations are "discursive" (which is to say, language-like); that reference is the only semantic

⁴⁷⁵ Fodor and Pylyshyn 2016: 128

property of mental or linguistic representations; that there are no such things as word meanings or conceptual contents... and so on'.⁴⁷⁶ I will describe these points in turn.

Reference depends on a causal chain between the referent and the tokening of the concept. This causal link determines the reference of both objects and their properties. To illustrate, suppose that I see a chair. There is a causal link between the chair and my forming a percept of it. The percept might then cause a token CHAIR₁ to be tokened. We then have a causal link between the concept CHAIR₁ and the chair. It is in virtue of this causal chain that CHAIR₁ refers to that chair. I might also have a more general concept of CHAIR that is causally linked to chairs that are outside of my own perceptual circle but that I may be causally connected to or that I am casually connected with in virtue of belonging to a community of speakers who are connected to those chairs.⁴⁷⁷ Alternatively, we may track the property of being a chair as it is caused by our experiencing chairs.

Fodor is also committed to the compositionality of thoughts. This is the idea that 'the content of a thought is entirely is entirely determined by its structure together with the content of its constituent concepts'.⁴⁷⁸ Call this the composition of content principle. For instance, if we have the thought LIONS ARE A TYPE OF MAMMAL, the truth conditions of this thought will be determined by the parts that make it up (LION, MAMMAL etc.) and the way that they are put together.⁴⁷⁹ The thought would not be true if mammals were a type of lion.

Fodor claims that 'Thoughts and concepts are individuated by their extensions *together with their vehicles*'.⁴⁸⁰ The first individuating feature is the referentialist claim, that concepts just have a reference as their content. So, the concept VENUS does not have an intension. For Fodor, a concept just has an extension and a vehicle. In this case, the extension is Venus. This leaves Fodor with a view

⁴⁷⁶ Ibid: 1

⁴⁷⁷ Ibid: 136, 138

⁴⁷⁸ Fodor 2010: 17

⁴⁷⁹ I am using caps to indicate the use of a concept rather than the use of natural language sentences.

⁴⁸⁰ Fodor and Pylyshyn 2016: 74

of propositions as sets of objects, as thoughts are composed out of concepts and have nothing else that could be contributed to their content.

However, VENUS is not identical with MORNING STAR, even though their extensions are identical.

This is often taken to suggest that meanings cannot just be extensions.⁴⁸¹ Fodor, however, explains the difference by appealing to the vehicles used to individuate the concepts. A vehicle is not something as elaborate as an intension. It is just that VENUS and MORNING STAR are tokens of different types. A concept is just a type, expressed in Mentalese, with an extension. 'Here... *syntax can do what senses were traditionally supposed to do*; that is, it can distinguish coextensive representations'.⁴⁸² The vehicles are different syntactic components of thoughts. If two thoughts differ from one another, it will be because they either have different extensions, or they include different vehicles. This is held to be true of both singular terms and predicates.⁴⁸³

Fodor also claims that "*Mental representations*" are the primitive bearers of intentional content'.⁴⁸⁴

This is the claim that mental representations do not need anything further for their content. They do not require anything else to have their content, such as a context of use. Whatever has content will have content in virtue of being in a relation to a mental representation. On this view a natural language sentence has its content in virtue of a mentalese sentence. This is a translational semantics '...which could be described in statements of the form "*abc*" means (=encodes) "*ijk*", where "*abc*" is a public-language form and "*ijk*" is Mentalese form' and mentalese sentences have content.⁴⁸⁵

Mentalese sentences have an extension independently of the context in which they are tokened because of Fodor's causal approach to reference, and the idea that reference is all there is to a concept's individuation (this leads to an atomistic view of concepts).⁴⁸⁶ This is because it is only

⁴⁸¹ Frege 1948: 210

⁴⁸² Fodor 2010: 61

⁴⁸³ Ibid: 199

⁴⁸⁴ Fodor 1998: 7

⁴⁸⁵ Carston 2002: 58

⁴⁸⁶ Fodor has other reasons for believing that concepts must be atomistic, see Fodor (1994) and (2000).

causal chains between a token of a concept and an object that determine reference. This excludes other contextual factors from playing a role in determining the referent. There are only referents for concepts. There is no other place for context to play a role in the thought. Context cannot, for instance, affect the intension of a thought because there is no intension. This does not mean that the context cannot play a role in guiding our translation from a natural language sentence to a Mentalese sentence. Only that the Mentalese sentence does not have its content affected by the context in a way that would make it context-sensitive. The referent will not change with the concept. Whilst the world is part of the mechanism of getting a referent, it does not remove underdeterminacy. The role that the context is playing here is at a stage of forming the concept, rather than at the purpose of applying or using that concept.

The extension of a concept remains fixed across contexts of use. So, for example, 'The concept DOG is locked to the property of being a dog and its extension is the set of (actual or possible) dogs'.⁴⁸⁷ Any change to this concept's extension would change the identity of that concept.

Recall also on this view that concepts are unlearned and, in that sense, innate.⁴⁸⁸ On this view, basic concepts are acquired because they are triggered by our experience. If we never encountered any doorknobs, that concept would never be triggered in us and thus never acquired. Non-basic concepts are made by conjoining two or more concepts to make a new one e.g. HORSE plus HORN for UNICORN.

Another point is that '...a thought can't be inexplicit with respect to its own content; there can't be more – or less – to a thought than there is to its content because a thought just *is* its content'.⁴⁸⁹ This rules out the possibility that there could be any more to a thought having the content that it does beyond the identity of that thought. That a thought cannot have content in addition to the content it

⁴⁸⁷ Fodor 2010: 141

⁴⁸⁸ Ibid:146

⁴⁸⁹ Fodor 2001: 14

has because of its composition also rules out some ways that context could play a role in determining content.

This point about the context not playing a role in determining the content of a thought needs some clarification. There are several ways in which composition might be understood. The context might play a role in determining the meaning, or the content of the constituents of a thought. One might want to say that it is possible that the context plays a role in determining the content of a concept before that concept is put into a syntactic whole. So, the concept GREEN might be modulated into GREEN*, where GREEN* might specify that the object be green only on the outside, and from there the sentence follows rules of composition. In this way the concept GREEN is altered into a new *ad hoc* concept that is intended to accommodate features of the context. It is claimed that '[t]he right model for an *ad hoc* concept seems to be a cluster of information all of which is occasion-specific and relevant for inference. Such a cluster, like a cluster of data points on a graph, has some clear members and outliers and there may be no sharp, non-arbitrary cut-off between these groups'.⁴⁹⁰ Whilst these concepts certainly reflect a pragmatic, context-based inference, and can be changed as needed, they still accommodate context-sensitivity in a way that can compete with Travis's view. Using modulation one can have both compositionality and context-sensitivity.⁴⁹¹

Fodor may be able to accept this form of context-sensitivity in thought. Whilst he rejects the possibility of a thought being inexplicit regarding its content, he does not seem to have a problem with using the basic concepts that we already have to create new non-basic ones.⁴⁹² What is important for Fodor's view is that thoughts and their concepts are not underdetermined. So that even if some additional work is needed to create that thought, once that thought is created its content is determinate. There is no further work for the context to do at the level of thought.

⁴⁹⁰ Allott and Texter 2012: 203

⁴⁹¹ Pagin and Westerståhl 2011: 19-20

⁴⁹² Fodor 2004: 85

Fodor also situates his understanding of representation in the CTM. Here ‘Accordingly, a mental process, qua computation, is a formal operation on syntactically structured mental representations’.⁴⁹³ Computations are only sensitive to the syntactic structure of what they act on, and do not require any understanding of the representation being operated on by the process. This is the recap of Fodor’s position. I will now include some additional details of Fodor’s view.

5.1.2 Additional Detail

I will now begin to add some additional details to Fodor’s view. Other important points that are worth emphasizing here include Fodor’s identification of Mentalese sentences with their content such that thoughts are the bedrock of representation. This is the crux of the disagreement between Travis and Fodor.

Fodor denies the compositionality of natural language. In so doing he accepts claim (a) of RRC. This serves as a premise in an argument that he makes, which goes as follows: ‘As between thought and language, whichever is compositional is the one that has content in the first instance. The evidence suggests strongly that language is not compositional. So, unless the evidence is misleading, it’s thought, rather than language, that has content in the first place’.⁴⁹⁴ So, natural language is context-sensitive. There may be a role for the context to play in determining the thought that an utterance expresses. There, the role that the context plays is as an aid in diagnosing what the speaker intended to convey, rather than playing a metaphysical role in determining what was conveyed.⁴⁹⁵ There is no room for context-sensitivity in determining the content of a thought.

Fodor takes thought’s purpose to be for the sake of ‘ascertaining truths’ rather than guiding our actions in our world.⁴⁹⁶ On Fodor’s view, thought’s paradigmatic role is in finding truth and he rejects

⁴⁹³ Fodor 2000: 11

⁴⁹⁴ Fodor 2001: 14

⁴⁹⁵ Fodor 2005: 104

⁴⁹⁶ Fodor 2010: 8

the idea that there needs to be any pragmatic element in thoughts.⁴⁹⁷ Thoughts need not have their content connected in any way to the purpose for which we employ those thoughts.

In summary, the key features of Fodor's view which Travis denies are:

- The content of a thought is a product of its components and their manner of composition, and nothing else.
- Mental representations are the primitive bearers of intentional content.
- Reference is achieved in virtue of a causal chain.
- Reference is the only semantic property of concepts.
- A thought cannot be inexplicit with respect to its content because a thought just is its content.

Crucially, these views commit Fodor to a denial of (b), the view that thoughts are necessarily context-sensitive. (b) is one of the distinctive claims made by Travis, so it will be important to understand why Travis thinks that (b) is true. In the next section I will present some of these reasons.

5.2 Travis's Criticisms

Travis is critical of Fodor's position. Fodor's position might be described as meaning deterministic. A key claim of these views is that it is possible for a series of concepts (or some other kind of representation) to determine the truth-conditions of a thought in which they appear independently of any further factors. In this case, the further factors would be the context in which the representation was tokened and any uses to which an agent would put the representation. For Fodor, thoughts are representations in this way. They determine their content without making use of the context. The mental state of tokening those concepts in that order is all there is to entertaining a content.

⁴⁹⁷ Ibid: 8

Travis's criticisms of this view take the form of rule-following arguments and are attempts to show that this kind of meaning determinism is false. Rules for applying concepts can be understood in a variety of ways. Travis most often takes rules to be something which determines the extension of a given concept, in this case they would be the function in intension or extension of that concept. The conclusion Travis draws is that there is no one right way of interpreting a rule so that it will always have one and only one extension. In other words, they are used to support a position according to which underdeterminacy is a universal feature of representations.⁴⁹⁸ If correct, this would rule out meaning deterministic views.

As Fodor's view seems to be an instance of meaning determinism, Fodor's position is deemed false by Travis. I will first give an account of meaning determinism and show that Fodor's position can be understood as a kind of meaning determinism. I will then present a version of the rule-following argument. This is important because accepting that thoughts could be context-sensitive won't be attractive otherwise. I will argue that the rule-following argument undermines Fodor's position. Though they do not guarantee that Travis's view is correct.

Meaning determinism is a set of assumptions. Kusch has a presentation of what he calls low brow meaning determinism. This goes as follows:

"Person x means Y by sign ' z '" is true if, and only if, x has a certain mental state (MS) that *constitutes* x 's meaning Y by " z ". Furthermore:

(MD-1) Immediate knowledge: x usually knows MS "immediately and with fair certainty".

(MD-2) Privacy: MS is an intrinsic state of x .

(MD-3) Grasping:

(MD-3.1) Grasping as cause: x 's act of grasping Y causes him – absent interference by other causes – to apply " z " in correct ways.

(MD-3.2) Grasping as intending: x 's act of grasping Y is tantamount to x 's forming intentions – or

⁴⁹⁸ A possible exception is made for mathematical propositions and some tautologies.

giving himself instructions – regarding a possibly infinite number of applications of “z” in the future.

(MD-3.3) Grasping as extrapolating: Since all learning sets are finite, grasping a meaning or concept has the character of an extrapolation.

(MD-3.4) Grasping as interpreting: Grasping a meaning or concept has the character of an interpretation.

(MD-3.5) Grasping as explanation: If x_1 and x_2 agree in all, or most, of their applications of “z”, the best explanation is that they have grasped the same concept Y .

(MD-4) Semantic normativity:

(MD-4.1) Non-blindness: In applying “z” on the basis of Y , x is not acting blindly.

(MD-4.2) Guidance: MS guides x on how to apply “z”.

(MD-4.3) Justification: x can justify his uses of “z” on the basis of MS .

(MD-4.4) Justification of unhesitating application: x can justify his unhesitating manner of applying “z” on the basis of MS .

(MD-4.5) Left -to-right interpretation of meaning conditionals: Meaning conditionals are to be read left to right.

(MD-5) Objectivity: x 's M[ental] S[tate] contains and determines (“in a queer way”) all future, potentially infinite, correct applications of “z”.

(MD-6) Classical realism: What gives a declarative sentence (DS) its meaning is the proposition it expresses. Propositions have truth-conditions. DS is *true* if, and only if, the proposition it expresses corresponds to a fact. Propositions are grasped; and conditions Md-1 to MD-5, as well as MD-7, apply to the act of grasping.

(MD-7) Metaphysical justification: The justification of our meaning sentences must come from ontological considerations'.⁴⁹⁹

⁴⁹⁹ Kusch 2006: 11- 12

Of relevance to Fodor's position is (MD-5). On the meaning deterministic view, the extension of a concept requires no additional pragmatic inference. The meaning is sufficient to determine the content of a representation. Furthermore, the content, or proposition, expressed is a metaphysical entity that we can use to determine the meaning of our sentences, whether they are sentences in natural language or Mentalese. So, on this view, there are mental states that have the same content across all contexts of use, and this is what justifies our deploying the representation as we do. Any differences in extension (keeping the world fixed) will be due to the use of different concepts or to their being composed in a different way.

On this view, it is the mental state of an individual that constitutes a case of a word representing what it does. This makes meaning deterministic views subscribers of translational semantics of natural language, according to which natural language sentences only get a semantics in virtue of being translated into a thought. So, whilst this description focuses on a sign "z" the relevant points here will be about the nature of the mental state, and that mental state's relationship to its content.

Fodor seems to take something like this view to be the correct one. He claims that the function of contextual information is not constitutive of the content of a thought. The context is not fed into thought via pragmatic inferences to determine the proposition that that thought expresses. Instead the context plays only a diagnostic role, to work out which possible interpretation is the correct one.⁵⁰⁰

This diagnostic role is only available whilst two people are communicating with one another, on the basis that it is absurd for someone to wonder what they had intended by a given thought. To illustrate this, consider a syntactically ambiguous sentence "I shot an elephant in my pyjamas".⁵⁰¹ This can mean that either the speaker was wearing pyjamas, or the elephant was. Whilst someone who hears this may use the context to infer which of the two was intended to be communicated, it is

⁵⁰⁰ Fodor 2005: 104

⁵⁰¹ Ibid: 107

not plausible that the speaker must infer which of the two possible propositions she intended to convey.

Fodor describes this as a distinction between metaphysical context-sensitivity of content, on which the content might in fact be underdetermined, and epistemic context-sensitivity of communication.⁵⁰² Whilst Fodor accepts the possibility of epistemic context-sensitivity of communication, he denies the possibility of metaphysical context-sensitivity of content or extension. So, it seems that Fodor cannot account for changes of extension by appealing to pragmatic inferences. He does not, however, believe that there are any such changes.

Fodor may want to deny that his view is really an instance of MD-5, as he takes the content of a thought to supervene on law-like connections between mental representations and things in the world.⁵⁰³ The connections are actual and possible causal relationships between objects and tokenings of those concepts.

This does not avoid it being the concept that determines its future instances. Whilst it is the causal interaction that causes the concept to be tokened, the concept type still seems to have its extension determined in advance. The concept type does the work because it is essential in making the process of tokenings occur in a way that is not random. If there was nothing intrinsic to the concept to determine which objects caused that concept to be tokened, then it is unclear why a given object should cause that concept to be tokened and not some other concept. It is the content of the concept that determines which causal interactions will token it. Furthermore, because the concept is individuated in part by its extension, if the extension of the concept were to change then it seems that the concept type would as well. That seems implausible. Alternatively, if you think that the vehicle is more important in determining the concept identity and can allow for the extension to change, then this seems to be a case of type-underdeterminacy instead. It is type-underdeterminacy

⁵⁰² Ibid: 106

⁵⁰³ Ibid: 109 - 110

because it would allow one concept to have several possible contents. Neither option is palatable for a Fodorian, so it seems that once the concept is acquired its extension will be determined by that concept.

There are several proper targets for the rule-following arguments. Most important here is MD-5, the claim that a mental state can determine the extension of a concept. On Fodor's view, a thought just is its content together with its vehicle and that is what it is to be that thought. So, the thought that is tokened will determine what the referents of that thought would be if it were true. Fodor takes it that the content of a concept is the set of all actual and possible things that fall under that concept. For instance, 'the content of CAT is the things belonging to the set of (actual or possible) cats'.⁵⁰⁴ The concept determines all its instances in virtue of these causal relations, which can be merely possible. The concept thereby determines all its referents. In which case, Fodor seems to ascribe to MD-5.

Travis's rule-following argument targets MD-5, as the content of a thought depends on more than just what the thought-vehicle is and on what the content and composition of its parts are. There remains a question of what the correspondence between a fact and a proposition should consist of. As a result, it seems that that there is some metaphysical context-sensitivity of content, not merely epistemic context-sensitivity of communication. At least, where subscribing to MD-5 is the alternative.

5.2.1 Rule-Following Arguments

The rule-following arguments are a collection of arguments that can be understood as targeting meaning deterministic views. I will present a rule-following argument as presented by Travis which questions in what way a thought can encompass certain states of affairs and not others. He describes this kind of rule-following argument as follows, where the fundamental relation is the relation of an object falling under a concept: 'The rule-following discussion concerns this fundamental relation. What is it for this [fundamental relation] to hold? When would it? What might answer these

⁵⁰⁴ Fodor and Pylyshyn 2015: 128

questions?'.⁵⁰⁵ The question is what is it that makes it the case that a given object, or a given input, can make the output true or false for a thought. Travis's arguments are intended to undermine the idea that a given concept can be to sufficient determine its extension.

One possibility is that the thought could express a function, either in intension or extension. A function in extension is a pairing of objects with truth values, and it is followed when one pairs the objects with the truth values as prescribed by the list. This is the understanding of a rule that Fodor seems to apply. A second understanding of a rule is as a function in intension that pairs possible inputs with an appropriate output based on a process that makes up the function.⁵⁰⁶ This function determines an output in terms of truth or falsity for each input, rather than listing the arguments and their values. On this approach, there does not need to be an exhaustive list for all the possible arguments. It can generate new values for new inputs as they are encountered. The claim would be that this covers the fundamental relation for future cases.

A thought just is its content, so just is its function, on Fodor's view. This view is described as follows: 'So no function could be a given one, F, if it mapped some argument into some value F did not, or if it had arguments F did not (e.g., if there had been objects there in fact are not). So, in the sense in which "is blue" names a certain function, it also names each argument of that function, or equally, each pair of an argument and the value that function assigns to it. Nothing could qualify as what that "is blue" named unless it took just those values for just those arguments'.⁵⁰⁷ Here, Travis is presenting one way one might identify and distinguish functions. The idea is that if you get different truth values for the same input, or different possible inputs, then there must be more than one function. So, a functions identity is tied to its pairings of arguments to truth values and there is no room for that function to include new arguments or to change the value assigned to those arguments. To do so is just to make a new function. This captures MD-5 well, in that this function

⁵⁰⁵ Travis 2011: 186

⁵⁰⁶ Church 1941: 3

⁵⁰⁷ Travis 2006: 36

determines its truth-value for all its arguments and if it did not then it would cease to be that function.

Fodor's view could allow for partial functions in extension, in which, for instance, clearly blue things are paired with BLUE and clearly not blue things are paired with NOT BLUE, leaving things which are neither clearly blue nor not blue unpaired. If so, Fodor's view can avoid metaphysical context-sensitivity if there is no evidence that this middle group can be allocated to either extension.

I will discuss an argument in favour of token-underdeterminacy where this is understood as:

Token-Underdeterminacy: A token of a structured representational item S is token-underdetermined if and only if for some possible states of affairs its truth-value is indeterminate (i.e. if and only if it determines a partial function from possible worlds to truth-values).⁵⁰⁸

If token-underdeterminacy is true of concepts, then the concept is not capable of determining what its extension should be in certain cases. The concept would not be able to determine its extension for all future cases, in which it should apply. If so, then there had better be a way of determining the extension, if we are to apply concepts to those future cases we encounter without a value attached to them. In either case the function does not determine its full extension. So, MD-5 will be false.

5.2.2 Token-Underdeterminacy

Token-underdetermined functions do not have extensions at all possible worlds. They only have extensions at some possible worlds. Travis seems to think that the functions we use are like this. So even when it is clear which of the possible functions we are using, that function may not inform us as to how it should be understood with respect to a given case. For a function in extension type account, it means that not all objects have a truth value paired with them.

⁵⁰⁸ Jaque 2017: 17-18. MacFarlane (2009: 233) has described this as context-sensitivity.

This is different from type-underdeterminacy. It is not a question of which function is the correct one, but what should the output of that function be given this (novel) input. In Travis's words 'it may be perfectly determinate which rule a given rule is, and still open for determination by occasions what following the rule in such-and-such case is to come to, or requires'.⁵⁰⁹ To give an example, suppose that there is an animal that appears to be a cat in all respects except that it also speaks Latin. In this case, it is unclear whether the animal is an instance of a CAT or not.

The target of this rule-following argument is a view on which the extension of a concept is determined independently of an occasion of use. Fodor's view, on which a concept's semantic value just is the set of referents determined by actual and possible causal relations, is such a view. Travis's position is that this requires the impossible of us. The extension cannot be determined in this way.

I will begin by building up Travis's arguments from cases of open texture. I will then present Travis's argument. The argument, roughly, rests on the idea that meaning determinism requires us to be able to think singular thoughts about objects which we are not in an epistemically rewarding (ER) relationship with.⁵¹⁰ This is because the concept of, for instance, CAT includes the extension of all cats on Fodor's view. But we cannot be in an ER with all instances of a concept. In which case, we are not able to think these thoughts. This argument is largely negative in that it attacks opposing views. I will present some responses to this argument, including a move back to an intensional account, before concluding in favour of token-underdeterminacy.

I will start this argument with the case of open texture. 'To say that a term is open textured is to point out that there are possible worlds where the application of this term is indefinite'.⁵¹¹ Most examples of this include cases in which there is something that is otherwise normal, but with some additional feature that makes it unclear whether we should count it as a member of the kind which it would otherwise be. Examples include gold that emits a new type of radiation, an animal with the

⁵⁰⁹ Travis 1989: 43

⁵¹⁰ Recanati 2012: 20

⁵¹¹ Margalit 1979: 141

body of a dog but the head of a human, and a cat that speaks Latin.⁵¹² In these cases it might well be unclear whether it really is gold, if it emits radiation, or whether it is a man whilst it has the body of a dog or a cat if it speaks Latin.

The point of these examples seems to be that ‘No set of rules can determine their [empirical terms] application for all possible situations’.⁵¹³ No matter how many scenarios we encounter in which we are able to decide whether “gold” applies, there may always be the possibility of something that is very gold like, but with some hitherto unencountered feature. These examples suggest that there are cases in which the content of a thought, the concepts involved, cannot determine all the possible extensions of a concept. The concept alone does not give an indication as to whether it can be correctly applied in these cases. This is the sense in which it is open.

On an open texture view it is possible that the content can still determine some of its own future extensions. These will be cases in which the content is not “open” to the object in that regard. So, gold that does not emit radiation or have some other strange feature will be gold without question even if we have not encountered it previously. These cases need not require any additional context to make an inference about the extension of that content. These facts are determined by the representation “gold” in these cases.

Travis, however, wants to adopt a more radical position. ‘Travis’s argument is in effect an argument for a radical form of open texture, in which it is underdetermined whether a thought is true or false or a concept applies or not *for any future application* of the thought or concept’.⁵¹⁴ It is not just that some potential objects are difficult to determine, but that the concepts themselves do not determine future extensions even when these cases are not so unusual. Even when a concept has been applied to an object on one occasion that does not mean that it can therefore be applied to that object

⁵¹² Margalit 1979

⁵¹³ Ibid: 141

⁵¹⁴ Hansen (unpublished)

again. This is because the occasion for applying the concept might vary. The correct application of a concept is occasion sensitive on this view. I will discuss the positive view in 5.3.

The problem which open texture presents for Fodor is not that there will be some inputs which do not get a truth value. This might just be an indication that the concept is indeterminate. The problem is that whilst the concept is indeterminate for these cases, it will not always be indeterminate. To illustrate, the water from a lake might appear blue when looking out over its surface, and yet be clear when in a bucket. In which case, water from that lake seems indeterminate as a candidate for blueness. The problem is that this indeterminacy seems to be subject to change. Sometimes that water will, reliably, cause BLUE to be tokened and sometimes it won't. These changes may be subject to various features, which I will discuss in section 5.3. The problem for Fodor's account is that causation alone does not have the resources to account for this change. On his view, '...there is no room for the idea that while speaking of that [being blue] (as the way my car, or, again, Lake Michigan, is) one *might* say any of various things'.⁵¹⁵ The water has the same causal properties that it always did and (I will argue) the concept BLUE has not changed either. Yet the truth value of THE WATER IS BLUE seems to have changed, even when considering the same water.

Travis also seems to assume that the extension of a concept cannot include objects that the subject is not in an epistemically rewarding (ER) relationship with. As there are many things which we cannot be in an ER relationship with we cannot include those things as members of the extension of a concept. One category of things that we cannot be in an ER relationship with are future things. In which case, we cannot include those objects as members of that set. This seems to be what Travis is getting at when he claims that a view on which the whole extension of a term is determined '...would make naming something presumably available to be named – such as things being under a bed, or being a grunter – require naming indefinitely many things then unavailable for being named'.⁵¹⁶ As

⁵¹⁵ Travis 2006: 35

⁵¹⁶ Ibid: 139

Fodor's view has the extension all actual and possible members of a concept's extension, future things must also be included in the extension. Travis's point is that this is not possible. We cannot think of those things in a singular way yet, so they cannot be included in the extension.

Fodor can deny that the extension of a concept can only include those things that a subject is in an ER relationship with, as he thinks that counterfactual causal chains from those objects to a tokening of that concept is sufficient. As counterfactual cases can include future cases, Fodor can just deny that we cannot be in the appropriate relationship to those members of the set. So future objects are still members of that set.

Furthermore, a Fodorian should deny that having a function in extension requires being able to think of the arguments that are paired with true in a singular way. He can claim that we do not need to think of the merely possible objects, that we cannot have encountered, in a singular way for them to be able to count as a member of that set (I take it that an ER relation to an object is necessary for singular thoughts about that object). All that matters is that they would so count. On such a view, the argument is paired with a truth value, but we are not able to think of the argument. Yet we would recognise it if we were to encounter it. The capacity to recognise these objects as members of a set seems to depend on a cognitive capacity to do so.

In response Travis can deny that counterfactual cases of causing a concept to be tokened are sufficient for that object to count as a member of the appropriate set. It is possible that a given input might sometimes cause the concept to be tokened and might not, as in the water in the lake example. Each tokening might be correct, such that the input does sometimes count as an instance of that concept and at other times does not. A causal interaction with the input alone is not sufficient to determine when that concept should be applied as the causal chain between the object and the subject can remain fixed whilst the extension varies. If there is variance, then Fodor's view is in trouble as his view is explicitly against this. This conclusion is largely driven by examples that Travis and others have considered, more on which below.

One such example revolves around the sentence “the leaves are green”. Suppose we have the thought that THE LEAVES ARE GREEN. Fodor’s view predicts that the truth-value of this thought will remain fixed for a given input (see 5.1.1). Travis challenges this with the THE LEAVES ARE GREEN case, which he presents as follows:

Pia’s Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning she reports, “That’s better. The leaves are green now.” She speaks the truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. “The leaves (on my tree) are green,” Pia says. “You can have those.” But now Pia speaks falsehood.⁵¹⁷

The point of this example is to take a given sentence and state of affairs and show that when the circumstances of evaluation changing, the extension of the sentence seems to change as well. The extension changes despite the input (the possible world under consideration) and content remaining constant. This can be used to suggest that the extension of a sentence is a context-sensitive affair.

Hansen has criticised this portrayal of the experiment, as the wording of the sentence changes in each case.⁵¹⁸ This is because the change in words leaves it open whether the difference in extension is due to a difference in the context or whether it is instead due to a difference in the expression. However, it seems that even when this and other factors are accounted for, people still react as Travis predicts.⁵¹⁹ This case can be modified so that Pia thinks in each situation THE LEAVES ARE GREEN. For the moment I am assuming that the output would be the same as in Travis’s original case.

It seems that in the first case what was thought was true and in the second what was thought was false.⁵²⁰ If this is the correct way to take this case, it is incompatible with Fodor’s view as he does not

⁵¹⁷ Travis 1997: 89

⁵¹⁸ Hansen 2018: 127

⁵¹⁹ Ibid: 129, 131

⁵²⁰ Travis 1989: 18-19

allow for variation in extension of a concept.⁵²¹ Fodor may want to say that one of these judgements of truth value is incorrect and that the concept is incorrectly applied.

Fodor takes incorrect applications of a concept to be ruled out on the basis of a causal asymmetry between correct and incorrect tokenings of a concept.⁵²² A fake horse (say, a well disguised zebra) will cause a token of HORSE only because the fake horse resembles an actual horse. HORSE, however, would still be tokened by horses even if horses did not resemble zebras, and even if zebras did not ever cause HORSE to be tokened.⁵²³ This asymmetry is intended to rule out cases of objects causing a concept to be tokened when those objects do not belong to that concept's extension. If one of the tokenings in the leaves case is incorrect, will we be able to tell which is incorrect based on asymmetry?

The asymmetry relation will not work for those cases in which an object might sometimes count as an instance of a concept and sometimes will not. There are not two objects that require resemblance in the leaves case, so there can be no asymmetry between the leaves and themselves. This way of accounting for variance in truth value will not work. So, when presented with that variation, Fodor's account needs a way of showing that there is no variance, as his view does not allow for this. In which case, one of the judgements that we have in Travis cases needs to be wrong. Either the thought is true in each context or false in both. But it seems that there is no principled way for this to happen on Fodor's view. He could say that what matters is whether the leaves appear green instead of being naturally green, or vice versa, but is not clear what reason he could give to motivate either option over the other one. Both have some plausibility and seem to be supported by people's

⁵²¹ There are numerous Travis cases, as they are called, that one might want to use in place of this, some of which I will discuss below. These include the sentence "the oven is hot", "there is meat on the carpet" and others. If one finds one of these Travis cases unconvincing then one can replace it with an alternative one. I will discuss and defend "the leaves are green" in more detail below. What matters is that one of these cases shows that at least one of these cases shows that the extension of a thought can vary for the same input.

⁵²² Fodor 1990: 91

⁵²³ Neander 2017: 416

intuitions on the matter. So, we should reject Fodor's view. I will discuss other ways in which one might resist Travis cases in section 5.3.1.2.⁵²⁴

This discussion began with the idea that a concept cannot include all its possible members as a part of their extension because that would require being able to think of all these possible arguments, which does not seem possible. Fodor's alternative was to rely on actual and counterfactual causal relationships between objects and tokenings of concepts. If correct, Fodor's position would not require that we are able to think of each of the possible arguments. However, Fodor's position is not able to account for cases in which the truth-value of a thought varies for the same input. Given that his view does not allow for context-sensitivity he needs a way to say that one of the judgements of truth value is incorrect. However, it seems that there is no easy way for a causal account to deal with this. This leaves the extensional account that Fodor wants to make use of in difficulty as it cannot account for this. In section 5.3.1.2 I will defend Travis cases from some common objections. The upshot of this discussion is that it is difficult for an extensional account such as Fodor's to allow for this variation. It also shows that an extensional account on which all the objects are determined in advance is implausible.

Instead of making use of an extensional account, Fodor could make use of an intensional account but that would be contrary to his naturalistic account (according to him). It appears that there is reason to believe that we cannot get a function that determines all the applications of that concept. If THE LEAVES ARE GREEN has the same content in each case, it still appears as though that function suffers from token-underdeterminacy. I will discuss some responses to this consideration as it applies to intensions in the next section.

⁵²⁴ Travis (2018: 309) claims that '...the ways for things to be of which we speak (and which we know to specify) are what are intrinsically susceptible to admitting of understandings'. So, it is intended that natural language show us something about the nature of concepts. I do not need to prove this is correct to prove that it is possible. See section 4.4.2 for discussion of analogy.

5.2.2.1 *Scepticism about Radical Open Texture*

Whilst one might accept that there are some cases of open texture, it might be denied that this can be a phenomenon that is as radical as Travis thinks it is. There may still be cases in which it seems that a given object definitely has a concept apply to it regardless of whether it is a future case or not and regardless of the circumstances. A case like this is presented by Stern who suggests that the extension is determinate regardless of whether we have encountered it or not. In particular, he examines whether DESK fails to determine its future cases.

‘...is it so clear that no extension is determined, regardless of understandings, by the very notion or meaning of “desk”? Suppose Jack had bought his desk, a dark cherry wood office model with drawers, locks and file, and [he] studies at it. Would the answer to the question whether Jack has a desk then “depend” on what the speaker means/understands by “desk”? No, if that is not a desk, it is hard to see what could be’.⁵²⁵

The idea is that whilst one might not have encountered the desk before, so that one cannot have a singular thought about it, it still falls under the extension of the concept DESK on all ways of applying DESK. That this object would (normally) be a desk seems uncontroversial, whether one had encountered it or not. So, it might be that the extension is not indeterminate because of having not encountered that desk. This can be cashed out in terms of their being an intension that determines the extension of at least prototypical cases of a concept in advance. Whilst this is not Fodor’s position, it is not Travis’s position either, and falls short of his radical token-underdeterminacy.

The challenge presented to the argument for token-underdeterminacy here is that we have a function in intension that does give a determinate output for some inputs. The desk case seems to

⁵²⁵ Stern 2003: 811. There is a slight misrepresentation of Travis’s view here, in that Stern takes the extension to depend on the speaker’s understanding of “desk”. This is not the case, as it seems that Travis (1989: 47) is concerned with what a reasonable judge would think counts as the extension of a term. The speaker may not be a reasonable judge in a particular case.

suggest that the output for that object will always be true when considering whether it is a desk and one need not have encountered the object in advance to know this. A function in intension can produce a determinate output for inputs if these inputs are encountered. So why should this intension be thought to be indeterminate? This is an important question. If we can accept an intensional account which does not admit of token-underdeterminacy then this may give an alternative way of preserving MD-5 and that understanding of the fundamental relation.

There are some reasons to think that things cannot be so simple as there being a determinate function in intension. One is that DESK does seem to exhibit open texture. For instance, a door over some crates might sometimes count as a desk but it won't always count as one.⁵²⁶ So, it is not obvious that token-underdeterminacy is avoided.

For another example, Pia might ask Sid for a red pen. On receiving a pen Pia might think THAT IS NOT A RED PEN. This may be for numerous reasons. Suppose that the pen contains red ink and is black on the outside. Pia's thought could be true for those cases in which she wanted a pen that had a red surface, or a pen that appeared red on yellow paper. Similarly, the thought will be false when she needed a pen with red ink.⁵²⁷

The point of this example is that this appears to be a thought which does leave its output indeterminate for certain inputs. Even if there are cases in which this might be determinate, for instance, if one had an object that was not at all a writing implement and gave a value false, there seem to be cases in which the output is left unclear. So, we have a case of token-underdeterminacy. One might object that this case can be resolved by changing the representation of the thought, to be something like THAT IS NOT A PEN WITH RED INK. However, even in this case it seems plausible to claim that the thought does not specify whether the ink should appear red on different colours of paper, or under different lights or when viewed through coloured lenses and so on for other

⁵²⁶ Travis 2002: 3-4

⁵²⁷ Jaque 2017: 96

variables that might affect the appearance of ink. Here it is not that we do not yet know whether those pens would count as red, and all that needs to happen is that we be acquainted with the object or circumstance, though that may be necessary. It is that the content does not determine what the output should be in this case.

So, whilst certain inputs do seem to give determinate values others do not, which suggests that there are cases of open texture. However, it does not suggest that open texture will be as radical as Travis makes it out to be. The concept might be determinate for some inputs, whilst yielding an indeterminate result for others. What emerges from this is that there are some concepts that exhibit open texture even when they are otherwise ordinary, so it cannot be assumed that concepts will be determinate.

This leads to two questions. One is to what extent might a given concept be open textured. A second is how many concepts may be open textured. Travis's position on both points is extreme. On the first he wants to claim that an open textured concept is token-underdetermined for all future inputs that it will have. For Travis, '[t]he concept as such admits of many applications, each excluding others. So it alone cannot assign an object, being as it is, a truth value'.⁵²⁸ This is the moral that Travis draws from his examples, and he also uses it as an argument to try to undermine the idea of a concept being a function from objects to truth values. Instead, it seems that there can only be a partial function, as a total function would be inconsistent. Even for those cases that are prototypes of their concept, if we have not encountered that object then we cannot have applied the concept to that object. For Travis, we could only be confident that if we were to encounter such an object then it would be correct to call it a desk.

On the second point Travis thinks that token-underdeterminacy will apply to general concepts, with a possible exception for cases like mathematics.⁵²⁹ General concepts are those that apply to at least

⁵²⁸ Travis 2011: 173

⁵²⁹ Travis 1989: 28

two objects and his examples often focus on predicates, so token-underdeterminacy extends at least that far. It may be that genuinely singular concepts are not subject to these kinds of concerns, as a singular thought is about its subject independently of any other predicates we might assign to that subject.⁵³⁰ Exceptions to this may be identity claims. More generally, Travis thinks that the conditions under which one can be said to have a singular thought vary with the occasion.⁵³¹ So, predicates will be token-underdetermined on this view.

For current purposes it is not a requirement that I prove both claims about open texture. Open texture, even for some cases, shows that MD-5 is false for at least these cases. The mental state that one is in (entertaining a given Mentalese sentence) does not determine all the correct applications of the concept. Travis's position can be understood as providing a means to show how that concept can be applied in those cases where a thought is token-underdetermined. If such a view works, then thoughts can be context-sensitive. In which case (b) is possible, if Travis's view does not appeal to thoughts which are context insensitive (or non-equivocal). I will present this view below and show that it does not appeal to such thoughts. This does not prove that (b) is true but makes it possible. In the rest of this chapter I will assume (b) is true and to see if it can avoid the absurdity the regress argument attributes to it. I will also use "underdeterminacy" to refer to token-underdeterminacy unless specified otherwise.

5.3 Travis's view

Travis denies the key components of Fodor's view as set out at the end of section 5.1. For Travis, the content of a thought is not determined by the content of its parts and the manner of its composition (in fact, he seems to maintain that the constituents of thoughts are determined by the thought and the occasion of its tokening).⁵³² Mental representations alone are not the primary bearers of intentional content, it is representations in conjunction with an act of judgement. Reference is not

⁵³⁰ Travis 2006: 64

⁵³¹ Ibid: 68

⁵³² Travis 2018

determined by a causal chain. Travis also makes room for there to be more than one semantic property. Travis allows for both content and meaning, though it is unclear what meaning is on his view.⁵³³ Finally, Travis also accepts that a thought can be inexplicit regarding its content. This occurs when the thought occurs without a purpose, which makes it impossible to judge what the extension of a given representation is. In this case, the thought has no clear content. I will give an account of each of these points. I will then give an overview of Travis's view.

A particularly hard problem in setting out this judgement-dependent view is to understand how it can establish a connection between a concept and its extension that does not make use of the rules that Travis believes do not determine extension whilst not making judgements arbitrary. In solving this problem Travis relies heavily on both a parochial sensibility and the purpose for which that representation was tokened. I will first spell out the points of contrast with Fodor's view and will then discuss judgement.

The first point is that content does not compose. Denying this form of the compositionality principle is standard fare for contextualist positions, including Travis's.⁵³⁴ Instead, the context plays a significant role in determining the content of a thought. Without a context 'Any representational form underdetermines when what has, or had, it would be true'.⁵³⁵ The scope of this claim includes natural language representations and Mentalese sentences. If the composition principle held, as Fodor suggests, then it would be possible for there to be representational forms which determined their truth conditions without the need for pragmatic inferences based on the context.

The second point follows from the first. Mental representations are not the primitive bearers of intentional content. If they were, and they did not compose, then it would not be clear what content they had. Their content would not be determinate, and this would seem to make the primary bearers

⁵³³ Travis 1997: 88

⁵³⁴ See, for instance, Jaque (2017b) and Carston (2002: 70-71). If they denied this, then it is unclear how they could maintain that space is left for the context to work.

⁵³⁵ Travis 2006: 32

of intentional content ambiguous. This would result in a considerable problem for the view, as it seems that we do in fact have mental states without this sort of ambiguity. For instance, if I think A BOY SAW A SPY HEADING TOWARDS THE SHOP, I know which of the possible contents that I am entertaining. Similarly, if I think I AM IN A BANK.

On Travis's view, the primary bearer of intentional content is a representation in conjunction with a judgement about that representation's extension. 'The main idea so far is that a semantic item, *W*, counts, on an occasion, *O*, as having semantics, *S* (or some particular semantic property *P*) just where all the other facts of *O* make it most reasonable to take *W* as having *S*, or to understand it as having, or in a way on which it does have *S*'.⁵³⁶ So it is a necessary condition for a representation to have the content that it does that it can be most reasonably judged to have that content. The thought alone is insufficient.

Similarly, the reference of a representation, such as *W*, is not determined by causal chain from an object, or members of a set of objects, to a concept. The extension on a given occasion is constrained by the meaning of the representation and made determinate by the reasonable judgement of someone who is informed about the occasion of use for that representation. The exact nature of this constraint is unclear, but it seems that the meaning will constrain the content so that it applies to those objects that the concept has been correctly applied to in the past. This need not be a consistent set of objects, as a concept or word can be applied to sometimes contradictory sets of objects.

Travis also maintains that there is more to a concept than its reference.⁵³⁷ There is also a partial function (a content) for concepts. This content cannot be something so substantial as to remove the role of the context in determining the extension. It must be weak enough to allow for alternative extensions. However, it cannot be so weak as to not rule out at least some potential extensions. If it

⁵³⁶ Travis 1989: 41

⁵³⁷ Travis 1997: 100

were, then it seems that we do not really have any fixed content at all. This is a worry that I will return to in section 5.4.

So, Travis opts for a judgement dependent view. On this view the extension of a concept is determined by judgements, provided that the judgement meets certain conditions. I will now give an account of judgement.

5.3.1 Judgement

Judgement is a technical term whose origins can be found in Frege. For Frege 'To make a judgement or to assert something is to recognise or to acknowledge the truth of a thought as the reference of the sentence that expresses the thought as its sense'.⁵³⁸ From this we get the idea that judgement is intimately connected to the truth-value of a proposition. When one decides that a proposition is true or false one has made a judgement about it. Judgements are necessary for connecting a content to a truth value.

This connection to truth remains at the core of Travis's notion of judging. Though for Travis more is required on the part of judgement to bring this connection between a content and its truth value about. What would count as an instance of the content's being true or false on a given occasion needs to be established. So, judgement will establish whether a given way the world is makes that content true or not. Given that there is no such thing as determinate content without a context, judgement is required to bring the context to bear on content. Ultimately, a thought cannot have a truth-value without a context. As a proposition is a bearer of truth or falsity, a proposition is composed of at least both a content and a judgement.

Judging itself is of central importance to Travis's version of Really Radical Contextualism (RRC). It is an act of judging that allows a thought to have an extension. What Travis means by judgement will

⁵³⁸ Carl 1994: 140

require unpacking and it is important that it is an *act* of judging that is required. To get a sense of its role in Travis's overall theory, the following quote will be of use.

Judging is engaging with the world precisely so as to be right or wrong about it according to how it is. A thought is the content of a judgement. It is, that is, a particular way of making one's fate – being right or wrong – depend on how the world is. It decides how the way things are *matters* to thus being right or wrong, how the world is to speak to that. It does that in fixing when things being as they are would be one's being right. The thought is that things are such that P; one is right just where things being as they are *is* things being such that P. The role of the thought is to fix when this would be.⁵³⁹

Before elaborating on this quote, it is worth noting how Travis uses the term "thought". In some places he uses it to mean something like the Mentalese sentence entertained, in line with how I have been using the term. However, in this passage the use is more in line with Frege's use of "thought". Frege says 'I call a thought something for which the question of truth arises. So I ascribe what is false to a thought just as much as what is true'.⁵⁴⁰ Here, "thought" is used to refer to a proposition, something that has the role of being true or false, and which need not be identical with any sentence.⁵⁴¹

However, Travis's concept of a thought as a proposition is not exactly Frege's. If it were, then it seems that Travis runs the risk of being targeted by his own rule-following argument. Whilst "thought" is sometimes taken to mean a proposition, and here we might understand a proposition as a function from possible worlds to truth-values, the nature of a proposition is distinct from the one that Frege had in mind. The key point is that contents are, on Travis's view, token-underdetermined. Their extension is not pre-determined for all possible worlds. Frege did not think that thoughts

⁵³⁹ Travis 2011: 232

⁵⁴⁰ Frege 1918/2010: 38

⁵⁴¹ Evidence for this reading of the passage also comes from other places where Travis (2011: 208) talks about judgement.

admitted underdeterminacy, so he could identify propositions with content.⁵⁴² Travis is not able to do this, so proposition will have to be a content plus something in addition. I will take propositions to be contents plus an act of judgement on Travis's view.

The token-underdeterminacy is removed, on an occasion, by judgement. The role of judgement is to decide whether a given state of affairs makes a given content true or not. Given a circumstance of evaluation and a content, a judgement is needed to determine whether the state of affairs makes that content true or not. This is necessary as the content is only a partial function and those unencountered cases do not yet get a value from the function alone.⁵⁴³

"Judgement" requires a few important features. These include the idea that we are pursuing truth when we make judgements of this kind, whether we judge something to be true or false is not felt to be an arbitrary matter. To perform these functions, judgements might make use of any number of contextual features. Finally, judgements are an activity that an agent engages in.

These features are summarised in the following quote. 'To judge that the salt is on the table is, per se, to see oneself as, given the world's impingement on him, having nothing else to say on the score: thinking otherwise would not be pursuing the goal, truth, hence not judging. We do not *choose* to judge'.⁵⁴⁴ Here, we get a sense of what judgements are meant to be. On the one hand, a judgement is a doing. It requires the agent to be in some sense actively engaging with the world. On the other, the agent cannot choose to judge however they like. They must attempt to judge to be in line with the way the world is.

So, an agent actively engages with the world but must use the information gathered in a reasonable way. Seeing that there is salt on the table and inferring that there is no salt would not be reasonable, for example. When judging, the conclusion needs to seem to be forced upon us by the evidence

⁵⁴² Frege 1918/2010: 38

⁵⁴³ It is worth noting that Travis does not seem to think that underdeterminacy will apply to all cases, he may make an exception for mathematical cases as these may not be type underdetermined (Travis 1989: 28). Singular terms may also be exempt.

⁵⁴⁴ Travis 2011: 5

available. In this case one cannot help but judge that the salt makes the thought THERE IS SALT ON THE TABLE true.⁵⁴⁵ For that circumstance of evaluation, including the purpose and other factors being considered, one cannot help but judge that *there is salt on the table* is true. This does not mean that they must judge correctly to be judging, but that they must be aiming to judge truly. Judgements can be incorrect and still be sincere cases of judging. To judge definitively one needs to judge as a reasonable judge would. An effective way of doing this is to be a reasonable judge.

This opens space for judging to be more or less involving. It may be only a very minor demand. For instance, seeing salt and judging that salt is on the table is active but easy for most agents to perform. By contrast, there may be more involving and difficult cases of judgement along the lines of what occurs in research. There may be a spectrum of cases of judgement. We sometimes make judgements in response to seeing the world or because of pursuing a deliberate line of enquiry.

Judgement is essential for resolving underdeterminacy. However, not any judgement will do. The judgement that X is a dog, for example, cannot be taken to be correct unless it is made by someone who knows what a dog is.⁵⁴⁶ Any other kind of judger's judgement would be insufficient to say correctly of X that it is dog. Instead, to be correct, it needs to be either a reasonable judge who makes the judgement or someone who judges as a reasonable judge would. Travis claims that 'you were correct in taking W to have the semantics you did take it to have if you thereby responded to all the other relevant facts as a reasonable judge would. To this we will add, "only if"'.⁵⁴⁷ This clause is necessary to keep concepts from being inappropriately applied. Someone who did not have a proper understanding of DOG might judge that non-dogs count as dogs. They might think of wolves as dogs,

⁵⁴⁵ Whilst Travis cases are always possible, in any given scenario the truth of a given proposition may seem unavoidable to someone. That you may judge differently in a different scenario does not mean that the judgement seems forced on you in another.

⁵⁴⁶ There will not be necessary and sufficient conditions for knowing what a dog is on Travis's view as the set of things that are dogs will not remain fixed on this view. So, stating what is involved in knowing what a dog is will not be an easy thing to do on this view, though it is not easy on any view. Knowing what a dog is will depend on knowing having encountered previous encounters with dogs, how to distinguish them from other animals and why we would want to classify something as a dog or not. It will also be partly a matter of how things seem to us as human beings, see section 5.3.3.

⁵⁴⁷ Travis 1989: 47

but this is to be mistaken as to the content of DOG, incorrectly judging what counts as an instance of DOG. At least, this would be the case if one was considering what sort of animal to own as a pet. (There may be cases in which a wolf would count as a DOG, if one was interested in the ability of an animal to produce offspring with one's pet).

5.3.1.1 Ideal Judges

The reasonable judges are not intended to be only ideals whose judgements ours are meant to conform to. It is also intended that there be actual reasonable judges. Evidence for this reading comes from the following quote. 'The hope would be that we may be able to put actual judges, beginning with ourselves, to substantive work as instruments for detecting what the reactions of a reasonable judge would be'.⁵⁴⁸ Furthermore, if some of us did not count as reasonable judges in some form, it seems as though we would have real difficulty in knowing what is a correct or incorrect applications of a concept (if one accepts the rest of Travis's view). Whether a possible world is paired with the truth-value true or not will depend on the reaction that a reasonable judge has.

In the absence of an actual reasonable judge we are left without access to this standard. In which case there still an ideal and correct and incorrect application of a concept. We just lack access to it. The role of actual reasonable judges is therefore to give us epistemic access to the standards of reasonable judges.

This is not to say that Travis should be committed to the view that we can always think truly. Just that there are cases in which we are knowledgeable enough that we should be able to apply these concepts reliably. Many of us are familiar enough with pencils that we should be able to recognise whether something is a pencil or not. There are some cases where we may not be able to, for instance, in poor lighting or a case in which someone has crafted pseudo-pencils. In the latter case, the standard in being a reasonable judge may be higher, as they need to be able to determine what is a pencil and what is a pseudo pencil. This does not seem to rule out the possibility of our being

⁵⁴⁸ Travis 1989: 48

reasonable judges. Rather, whether one can be a reasonable judge will depend on the circumstances in which one makes the judgement. Similarly, being a reasonable judge in one area does not make one a reasonable judge in another area. There are several factors that a reasonable judge needs to be sensitive to, which I will discuss in 5.3.2 and 5.3.3.

5.3.1.2 The Leaves are Green

Judgement also allows for Travis to bridge the gap between concepts and their extension in a way that allows the context to play a role in setting the extension. A reasonable judge does not just have to consider the concept and the extension, but also the occasion and purpose for making the judgement. As such, reasonable judges are meant to take the most reasonable interpretation of a thought. It is also because we consider what is reasonable that there is room for meaning to be an occasion sensitive affair.⁵⁴⁹

Occasion-sensitivity is illustrated by the thought THE LEAVES ARE GREEN. To recap, Pia's Japanese maple is full of russet leaves. She paints them green and thinks THE LEAVES ARE GREEN. She thinks the truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. Pia thinks THE LEAVES ARE GREEN. But now Pia's thought seems false.⁵⁵⁰ Here, the extension changes despite the input (the possible world under consideration) and content remaining constant. This suggests that the extension of a sentence is a context-sensitive affair.

It is important that Travis argues against other interpretations of these experiments. One important alternative is that GREEN is a vague concept. If this were the case then we would expect the cases that Travis talks about to be borderline cases of being green, which we might sometimes take to be green, and sometimes not. This does not seem like a correct description of what happens here. It is not that the painted leaves are on the periphery of greenness. It's just that they are sometimes what

⁵⁴⁹ Travis 2006: 128

⁵⁵⁰ Travis 1997: 89

we would count as green and sometimes not. In each case they are firmly “in the middle” or not of greenness, not on the periphery.

A second alternative is that there are several concepts of green, say GREEN and GREEN*. In the first case, the leaves are GREEN in that they have a green surface. In the second case, they are not GREEN*, which may be something like not biologically green. This alternative explanation here is that there may be *Ad Hoc* concepts that we compose as needed to suit the demands of the situation.⁵⁵¹ If so, the correct description of the case is one in which the concept is not constant in each case. So, when Pia has painted the leaves, she may use a concept GREEN and when she thinks about giving the leaves to her friend, she uses a different concept GREEN*.

If this is the correct interpretation of what happens then we do not have a case of a single thought having different truth values in different contexts. Rather, we have the construction of new concepts leading to two distinct thoughts with two distinct truth values. The result would be two distinct mentalese sentences with different truth conditions and therefore different truth-values for the same input. This is compatible with Fodor’s view, minus the intensions.

However, if the arguments in favour of underdeterminacy are convincing (as I am assuming) then this interpretation cannot be all that there is to that case. Whilst there may be two mentalese sentences, that they have the content that they do will not simply be a matter of their having distinct concepts GREEN and GREEN*. There will have to be a judgement that applies the concept to that possible world, or not. To illustrate that the content might still be token-underdetermined, we can take GREEN to be painted green and GREEN* to be naturally green. Though in each case we might want to distinguish different potential contents that each could be judged to express. For instance, painted

⁵⁵¹ There are several ways one might understand the production of *ad hoc* concepts. One is that the extension of a concept is altered (Wilson and Carston 2007). Another is that the intension is changed (Allott and Textor 2012). Either option could be performed by conjoining different concepts together. If there was a thin concept GREEN_IN_SOME_WAY that was conjoined with a more determinate way of being green e.g. PAINTED that would still be an *ad hoc* concept.

green might not always apply if the leaves have only been stippled with paint or have dots on them. So, the need for someone to apply the concept recurs.

Underdeterminacy might be difficult to establish if one thinks that the concepts in this case are *ad hoc* and were made to pick out a property in the world then the truth-value may well be determined for this case. The concept is like a singular thought in that it was made to pick out *that* property. It is significant that on such a view the *ad hoc* concepts would not be singular thoughts, they are still general thoughts as they are a concept that plays the role of a predicate and these are properties that might be instanced by more than one item. That these *ad hoc* concepts are custom built for a given task does not entail that the contents associated with both are more than partial functions. There will still be unencountered cases and occasions without a truth-value parried to them. Furthermore, it is not obvious that this is the best way to understand the formation of *ad hoc* concepts. It does not seem that the concepts are formed specifically to fit that state of affairs or not. Instead, it seems that they are composed to make the distinction between painted green and naturally green which fits the purposes of the thinker. In which case they will be open textured.

Travis concludes that GREEN does not have a fixed extension and requires a judgement to determine the extension on each occasion. The leaves case illustrates that judgement should consider features of the context. In this case, what the leaves are intended to be used for. Where the leaves are intended for use in botany it is unreasonable not to take this into account, and similarly for other features of the context such as the leaves' natural colour.

5.3.1.3 Comparison to Wright

This suggests that the criteria for being a reasonable judge are subject to the conditions in which one is judging. If one is going to judge successfully then certain conditions will have to be met. A similar view is suggested by Wright, according to which one can only give extension determining judgements under optimal conditions, called C-conditions. Wright claims that certain kinds of judgement cannot be wrong when made under C-conditions, for instance, one cannot be wrong about whether

something appears blue to you or not.⁵⁵² For Wright, these judgements cannot be wrong because these judgements determine the extension of these concepts. The judgements are extension determining.⁵⁵³ Wright claims that semantic properties are also determined by our best judgements made under C-conditions.⁵⁵⁴

Travis's view is like Wright's in that to be a reasonable judge one needs to be making the judgement under suitable conditions and that it is not possible for these judges to be mistaken when these conditions hold. These judges are taken to set the standard for the application of the concepts that they qualify as reasonable judges for. So, Travis's view is like Wright's in allowing a kind of judgement to determine the extension of a concept.

However, the two views are distinct. There are differences in the conditions under which one can make an extension determining judgement. For Wright, the C-conditions for judging whether something is blue are that '...the surface must be in full view and in good light, relatively stationary, and not too far away; and the thinker must know which object is in question, must observe it attentively, must be possessed of normal visual equipment and be otherwise cognitively lucid, and must be competent with the concept *blue*. In addition, the thinker must be free of doubt about the satisfaction of any of these conditions...'.⁵⁵⁵ These conditions do not make it possible for the judgement to determine the extension differently on different occasions.

Recall Travis's take on "The leaves are green" which is true when Pia has just painted them but false when she tries to give them to her botanist friend. Here, Pia may be under the same C-conditions. Yet in one case she has made an error and in the other she has not. This is because reasonable judgement has different conditions depending on the purpose for which one is making the judgement. For instance, if Pia's friend does not need green leaves for an experiment, but rather

⁵⁵² Wright 1989: 130-131

⁵⁵³ Wright 2001: 192

⁵⁵⁴ Ibid: 212

⁵⁵⁵ Ibid: 192-193

needs leaves that appear green under a neon light, then it will not matter how the leaves appear under optimal or normal lighting conditions.

This helps to illustrate how radical a position Travis's is. The conditions for being a reasonable judge vary with the occasion. To elaborate '...our capacity to recognize what so counts, and what not, is not just a capacity to go by those marks [having a snout] —a preparedness so to respond—blind to anything further which might bear (or not) on whether an object counted as a pig. In that case such would simply be how we respond to those marks; such is what we do. It would not be sensitivity to reason'.⁵⁵⁶ The idea being that we should take into consideration more than just whether the C-conditions are met but how the circumstances in which we are making a judgement should effect that judgement. This includes the purpose for which that judgement is being made and which distinctions we are hoping to draw. The conditions for reasonable judgement are not fixed across occasions of use, so the extension of a concept is liable to vary across tokenings. A key point for Travis is that understanding the purpose for which a concept is to be applied is essential for the judgement about the extension of that concept to be extension determining. I will elaborate on this in section 5.3.2.

Where there is underdeterminacy we find ourselves in a situation in which we have an object or occasion for judging that we have not yet encountered for that function and so do not know what the truth-value should be. So, we are not yet sure whether the world in which we find ourselves would give the value true or false for this function. In this case the judgement is to decide whether this new input should give the value true or false.

Similarly, for Travis, a proposition is what one gets when making a judgement. 'The proposition is then identified by what would be, recognizably, *someone's* being exposed to risk in a particular way'.⁵⁵⁷ Here the risk might be something minimal, such as the risk of believing something that is

⁵⁵⁶ Travis 2018: 82

⁵⁵⁷ Travis 2011: 209

false and thereby failing as an epistemic agent. It may also be something more substantial than that, such as the risk that one will miss an important meeting, be bitten by a snake that one had mistaken for rope or something similar. This conception of a proposition is related to the idea that generating an extension requires a purpose.

Furthermore, Travis's response to the rule-following arguments is unlike other responses to the rule-following argument in that it does not give up on the idea of there being truth-conditions. This makes it unlike the Kripkenstein response where truth conditions for natural language are abandoned in favour of assertability conditions.⁵⁵⁸ However, it is unlike other positions that maintain that we can speak about truth in its radical reinterpretation of propositions. A content is not able to determine its extension for future cases. When a proposition is true is not determined by the facts of the world and the nature of the content alone. Propositions require judgement, one case at a time.

5.3.1.4 Judgement in the Mental Life of Agent

One question worth asking here is what a judgement looks like in the mental life of an agent. So far, it seems to be a kind of propositional attitude that one can take to a given proposition, of the form S judges that P is true in these circumstances. Where S is the subject and P is a proposition. This seems to fit some of the features attributed to judgement. For instance, that it be in pursuit of the truth, and that it can be prompted by witnessing the world to be forced to conclude that P is true. It also makes sense of the idea that a proposition is the content of a judgement.

One might worry that if RRC is true then the propositional attitudes we entertain in making judgements are also underdetermined. For Travis, there seems to be no way of denying that this is so, given how pervasive he takes underdeterminacy to be. In judging that a given proposition is true

⁵⁵⁸ Kripke 1982: 74. Dobler 2017 suggests a similar reading of Travis according to which there are constant truth-conditions for given utterances but varying acceptability conditions, see chapter 1. This differs from Kripke in that it accepts that there are still truth conditions but differs from Travis in claiming that there are context insensitive contents.

on this occasion, that judgement may itself be underdetermined and so not apply determinately in that case or in future cases.

The way out of this is to emphasize the relationship between judgement and action. Judgements are made to help us to navigate the world in which we live in. For instance, being able to make judgements about what is and isn't a snake, what is a dangerous situation and what a good opportunity is are useful skills. When we make a judgement about a representation that judgement needs to have a purpose, or it will have no connection to action. When a judgement gets made it ties our actions to the world in a particular way. So, Travis claims that 'action is the foundation of representation' and that '[w]hat we judge is what can matter to what to do – how to treat things'.⁵⁵⁹ The reason that these judgements do not need any further judgement to make them propositional is that they are able to play their role in guiding our behaviour as they are.

In judging that something counts as a rattlesnake one has decided how to treat it, for instance, keeping a safe distance. For that action, the future cases do not matter. That we can act on it makes it determinate enough for that case. That a given judgement allows us to act in a correct way does not mean that the judgement was correct, but it does help to tie the representation to the world. So, even an incorrect judgement might still allow us to treat situations in the correct way e.g. mistaking a rattlesnake for a cottonmouth snake will still allow one to avoid it etc. What matters in this case, let's suppose, is that both species are dangerous. This will not always be the case and is not the only factor that matters. I will discuss other criteria that reasonable judges make use of below.

This means that judgements are not meant to resolve all the possible underdeterminacy for a given thought. But for Travis that is acceptable. All that a successful judgement needs to do is to decide for

⁵⁵⁹ Travis 2011: 209. When it comes to co-extensional concepts a reasonable judge will judge them to be so. The concepts can still be distinguished in virtue of their having different partial functions that happen to be co-extensional e.g. trilaterals and triangles. As in Fodor's case these concepts may also be distinguished by their vehicles.

this occasion whether the output should be true or false. If it does that much then we know how we should act on that occasion, or at least, we know something that can help us to decide how to act.

To sum up, the role of judgement is to bring the world and representations together so that it is clear when contents can be applied. This application of representations is what allows them to have truth-values. In judging a thought to be true, we should not be in a position in which we feel we can choose whether it is true or not. Instead, it must be as though the truth-value is forced upon us by the evidence. I have briefly tried to flesh out this view of judgement dependence and compare the view to Wright's view of extension determining judgements. An important difference that emerged here was the role of the purpose for applying a given concept plays in Travis's view. The purpose seems to play a role in determining when one can be a reasonable judge.

In the next section I will present purpose and its role in making judgements. I will then present the second condition that Travis places on judgement, which is that there be a human parochial sensibility from which to make judgements. Without these features one is not able to make reasonable judgements, regardless of whether one would otherwise count as a reasonable judge. These play a role similar to C-conditions on Wright's account. This is not an exhaustive list of the pre-conditions that one needs to count as a reasonable judge. For instance, a reasonable judge on topic X will need to have a good amount of knowledge of topic X.

5.3.2 Precondition 1 – Purpose.

Travis is committed to the claim that purpose is a necessary condition for representation (I will take purpose of a representation to be interchangeable with its point). For there to be a representation which expresses content, there must be a purpose in tokening that representation. 'If the driving idea here were to be put into a slogan, it might be this: Content is inseparable from point. What is communicated in *our* words lies, inseparably, in what we would expect of them. How our words represent things is a matter, and not detachable from, their (recognisable) import for our lives'.⁵⁶⁰ It

⁵⁶⁰ Travis 2006: 33

is important to note that purpose is not taken to be sufficient for content. Rather, purpose is necessary to resolve underdeterminacy by determining the truth-value of that content in that case. The aim of this section is to elaborate on how purpose might do this.

Whilst Travis focuses on natural language, in the above quote, it is important that his point applies more generally. Thoughts also have a content, and to have a content requires that they have a purpose. The role of purpose is to help to fix the appropriate standard for truth on that occasion.⁵⁶¹

Travis has a useful illustration. 'One would not count a 140°C oven as hot for purposes of baking pizza. A hot oven, for that purpose, is much hotter. By contrast, for purposes of removing a rack bare-handed a 140° oven would count as hot'.⁵⁶² This shows how the extension might change with the purpose of that representation. If I am wondering whether I need to use oven gloves or not, then the truth conditions of the thought THE OVEN IS HOT will be different from a case in which I am wondering whether to put a pizza in the oven.

Travis also claims that 'It would be useless to be told that the oven is hot if we had no idea of the standard by which that was to be judged. Our perceptions of purposes to be served, and uses words would have in serving them, provide just such a standard.'⁵⁶³ Here the point is that there is no such thing as simply being hot. That is not something that the concept, or the word, is able to do. Here having a sense of purpose is to understand what actions one might need to take and then to understand the factors that would be relevant to determining this.

That the purpose for which we employ a representation can change when a given representation counts as true or not means that purpose is also playing a deeper role in representation. The conditions under which one can be said to make an extension determining judgement vary with the occasion of making that judgement. This suggests that purpose not only changes the relevant

⁵⁶¹ Travis 2008: 104

⁵⁶² Ibid: 103

⁵⁶³ Ibid: 104

standard which an object needs to meet to count as, for instance, hot. It also changes the conditions under which one is a reasonable judge. To judge that 140°C is hot for a human to touch without protection requires knowing that humans will be burned by things of that temperature whereas judging that 140°C is not hot for the purpose of cooking a pizza requires knowing what the usual temperatures are for cooking pizza. So not only does purpose set a standard for what is to count as something that would make a given thought true, purpose also changes the conditions under which one can be said to count as a reasonable judge.

There are some extant and well-rehearsed challenges to Travis's claim that purpose is necessary for representation (e.g. see Fodor and Pylyshyn 2015: 60, Cappelen & Lepore 2005, Borg 2012: 73, 2004: 236-7, 2016 etc). Unfortunately, however, exploring these challenges would take us too far from the central concerns of this thesis. Thus, in what follows I will simply take it that Travis is right to maintain that purpose is necessary for representation, using it as a working hypothesis to explore his view and the possibilities it opens for context-sensitive thought.

5.3.2.1 Clarifying Purpose

I will now clarify the role of purpose. It is important to clarify what purpose is intended to accomplish and how it might do so, if Travis's position that purpose is necessary for content is to be assessed. I will suggest that the role of purpose in making judgements is to make it apparent to us as judges what is at stake in a given scenario. This, in turn, can direct us to the ways in which a given concept is open textured and to which of these openings is the relevant one to be closed for the current judgement. This should also give us a good idea of what we will need to know to judge reasonably. The idea is not that judgement should resolve all the underdeterminacy of a given concept. It need only resolve enough for the current purposes. Once that judgement is made, it does not follow that the same input will always yield that truth values in future cases, as the purposes for applying that concept may be different in those cases. Open texture is only closed for a given occasion and not in general.

In clarifying purpose's role in judgement, it will be useful to turn to some examples. One example is the thought that "There is red meat on the white rug".⁵⁶⁴ In this example as Sid walks across a room '...the kidneys he has bought for the mixed grill fall in their butcher paper to the rug'.⁵⁶⁵ Suppose that this is a case of underdeterminacy. The truth of the thought that THERE IS MEAT ON THE RUG may vary depending on the purpose one has in mind on that occasion.

If you're wondering whether you now need to clean the carpet then you may think that it is false, providing that the butcher's paper is thick enough to keep the carpet clean. If your purpose is to decide whether to pick it up, then it may count as true. This will also be a case in which one is prepared to admit that kidneys are meat, rather than offal. This may also depend on one's purpose. For example, if the kidneys are to play the role of meat in a meal then they may count as meat. Otherwise they may not, and the thought will be false.

This example goes some of the way towards showing how truth may vary with the purpose one has in mind. Many of Travis's examples seem to direct one to this conclusion. However, these examples also illustrate why purpose is necessary to establishing an extension. There are several possible outputs that one could have for this partial function. What's more is that some of these conflicting extensions seem reasonable to us. However, we cannot make use of all them at once. So, a decision needs to be made, and this decision cannot be made based solely on the meaning of the concepts and the way that the world is. So, the function of purpose is to guide our judgements in deciding how to extend a given function. It may serve as an input to a partial function, alongside a possible world. In this way, purpose helps us to conceive of the world in a way that allows us to pursue our goals effectively.

It is worth noting that Travis does not take judgements to be solely determined by purpose. For Travis, judgements must also have the aim of being true. However, purpose is required to keep this

⁵⁶⁴ Travis 2011: 243

⁵⁶⁵ Ibid: 243

grounded in our affairs and ensure that we have the tools required to make judgements. This means that concepts are kept from becoming too detached from objects and actions in the world. Purpose means that when we judge mistakenly there are consequences that may frustrate this purpose. For instance, a mistaken judgement about the departure time of a train may frustrate one's travel plans.

More fundamentally to judging, purpose provides a way of helping us to decide which factors are relevant for the assessment of the world in relation to a proposition. Purpose helps us to decide what counts as something being a certain way and what that way might be by making it clear why it is important to classify things in a given way. It does so by directing us to the relevant opening in the open texture of the concept. So, whilst purpose is necessary, it is not sufficient for establishing an extension. Whilst it can guide us towards the relevant opening in a concept, it does not give us the capacity to see when something is fit for that purpose or not.⁵⁶⁶

This means that there are cases in which purpose might be overridden by the need to have a precise understanding of the world. So, in the case of thinking *THERE IS MEAT ON THE CARPET* when it is a kidney on the floor, one might judge this false if one cares more about being precise than whether something needs to be picked up off the carpet. Priorities might vary within the same individual across time and sometimes rapidly. So, a person might judge that *there is meat on the carpet* when considering the offal. They may then realise that kidneys are not actually meat. Nevertheless, they still pick them up. Here their purpose has not changed but their judgement about what is on the floor

⁵⁶⁶ Travis's view has been compared to some other views, such as MacFarlane's "counts-as" view. On MacFarlane's (2007:246) view, the circumstance of evaluation consists of both a possible world and a "counts-as" parameter. The counts-as parameter serves to fix what counts as being e.g. tall in that circumstance. This parameter determines the extension in that scenario. From an exegetical point of view, this is not a way of interpreting Travis's position. For instance, Travis (1997: 93) claims that '...perhaps for any set of parameters, further possible factors would yield more than one distinguishable thing to be said for fixed values of those'. When discussing sentences Travis (2006a: 41) denies that the proposition expressed by a sentence is '...predictable (as a function from some set of parameters)'. However, this may be more of a difference of emphasis than of content. It may be that the counts-as parameter is itself determined in a way that Travis suggests, this an option pursued by Collins and Dobler (2018: 7-8). That is, it is sensitive to the purpose of the thought and the parochial perspective of the thinker. In which case, we still have underdeterminacy and there may still be a way of understanding the counts-as parameter as judgement dependent in the way that Travis describes. In which case, one can describe the position as one likes.

has. So, whilst purpose is important, it is not the be all and end all of judgement. Purpose may have an effect on how precise one's judgements need to be. It may sometimes be important that judgements be made quickly or frugally and at other times the level of precision may be most important. Part of the task of a reasonable judge is to balance these two factors appropriately, but both are necessary to any reasonable judgement.

In this section I have discussed purpose as a necessary condition for making judgements that can resolve token-underdeterminacy. Purpose is necessary to direct judgements towards the factors that might be relevant for determining a given judgement on an occasion. Purpose also helps to ground our judgements in the world and create a cost to judging poorly. I have also tried to defend this approach from objections. In the next section I will consider the second key precondition for extension determining judgements that Travis considers. This is the parochial sensibility that allows us to recognise when a given state of affairs might be useful for our purposes.

5.3.3 Precondition 2 – The Parochial sensibility.

For Travis there is a uniquely human parochial form of thought which he describes as follows:

I will refer to as “parochial” any form, or shape, of some’s being thought which is not required simply by the demand of being a thinker as such, so that there is room for there to be *thinkers* whose thought lacked that feature – so to any form of thought for which there might be others.⁵⁶⁷

The idea that we operate with a parochial form of thought is a central assumption of Travis’s view. In particular, he claims ‘[w]e can make no sense of the idea of a way for things to be, identified as the way it is independent of any special parochial sense for novel understandings as to which way it is to be, which requires, merely in being what it is, some unique range of novel understandings’.⁵⁶⁸ Here Travis is making the claim that we cannot have any clear idea of the extension of a concept where

⁵⁶⁷ Travis 2011: 1

⁵⁶⁸ Travis 2006: 138

this extension extends to all of the states of affairs that would make the thought true. Furthermore, it shows that Travis takes a parochial sensibility (or perspective) to be necessary for determining extension. The question for this section is what the parochial sensibility does which makes it a necessary condition on judgement.

Purpose helps to set the standard by which we should make a judgement on an occasion and thereby determine the conditions for making an extension determining judgement. However, recognising what counts as something meeting this standard is what the parochial sensibility is meant to allow for. Purpose points us in the relevant direction, but the parochial sensibility is needed to recognise what counts (it is also likely that our parochial sensibility plays a role in determining what our purposes are likely to be e.g. obtaining shelter and so on are likely to be important because we generally see this as valuable). Travis's claim is that 'for a rule to count as saying to do *this* (identified by doing it) is for it to say that on a certain parochial understanding of it; to do so on a given thinker's way of understanding it (on what *they* are prepared to recognise as the right way)'.⁵⁶⁹ This suggests that it is the perspective that we have on the world that allows us to decide whether something is fit for a given purpose. This parochial sensibility is a capacity to recognise when an object can be counted as falling under a given concept for that purpose. It is a capacity to recognise when something is useful for a given purpose. A part of this perspective will be that certain things are generally considered important or relevant to us as human beings e.g. food and shelter.

The role that parochial sensibilities play in our judging runs deep, as hinted at in the following quote.

'The conceptual is autonomous. When we think something to be so, we connect to some bit of it which reaches just where it does, independent of *anything* else (except, perhaps, that there be those cases to reach to). It is just that the parochial sensibility is needed to forge any such connection'.⁵⁷⁰

The conceptual is Travis's term for the set of concepts that we have available to us. Importantly,

⁵⁶⁹ Ibid: 115

⁵⁷⁰ Travis 2011: 23

these concepts are abstract, and need to be applied to concrete individuals or objects. The fundamental relation is '...for a generality to reach all the way to a particular case'.⁵⁷¹ Elsewhere, Travis puts the point like this '... the parochial gives Frege's fundamental logical relation its purchase: allows for facts as which ways for things to be are instanced by the ways things are, and which are instances of given ways there are for things to be'.⁵⁷² The parochial sensibility helps us to judge whether a given object should count as an instance of a given concept. This is an important point. The claim is that for us to be able to apply concepts to things in the world, to be able make judgements, we must be able to make use of our unique way of thinking as a species.

Following on from the rule-following argument Travis claims that there are no interpretation-proof rules, for a rule to count as saying to do x is for it to do so on a certain parochial understanding of it.⁵⁷³ He also claims that our parochial sensibilities contribute to fixing what it is that is to be true.⁵⁷⁴ However, the point is not just that the parochial sensibility plays a role in our judgement simply because we cannot help but use it. Without a parochial sensibility there can be no such thing as an object instancing a (general) concept. An object can only count as an instance of a given concept if it appears to be an instance of that concept from a given perspective. So, a reasonable judge must have a perspective.

This raises questions. Presumably, the parochial sensibility here is playing a role in our judgements, which is what determines the instances of a concept. But what role does it perform that was not performed by the purpose of instancing the concept. In this case, the parochial sensibility is intended to be that which gives us the capacity to recognise a way the world is as an instance of a given concept. The purpose serves to guide the parochial sensibility on this matter. Given that the concepts we have available to us do not have a fixed extension and cannot provide us with the extension of a concept (as per the rule-following argument) we require a way of fixing the extension on each

⁵⁷¹ Ibid: 186

⁵⁷² Ibid: 221

⁵⁷³ Travis 2006: 115

⁵⁷⁴ Ibid: 118

occasion. It is the role of the parochial sensibility to provide us with a means of recognising when a state of affairs counts as an instance of a proposition being true or not. Where the purpose provides a standard, the parochial sensibility is what allows us to understand when this standard has been met. It is what provides us with a means of understanding how concepts are to be related to the world.⁵⁷⁵

How does the parochial sensibility do this? It seems that the parochial sensibility provides us with cognitive capacities, and it is these capacities that allow us to apply concepts to the world.⁵⁷⁶ If we had no such parochial sensibility, we would be left with concepts which we could not apply. The parochial sensibility itself, and the capacity to apply concepts as we do, may simply be a result of our neurology and physiology. In this respect, Travis's position has a point of contact with Fodor's. Travis also seems to require there to be something in human nature that allows us to apply concepts to things in the world. Where Fodor claims that doorknobs just are the sorts of things that make minds like ours jump to DOORKNOB, Travis seems to claim that minds like ours can apply DOORKNOB to doorknobs. We have the capacity to recognise when this is a correct application of a concept. Taking this route puts Travis and Fodor on equal ground in this respect. Both are dependent, at some point, on our cognition working the way it does because of some biological facts about us as humans. A difference is that other factors may shape our parochial sensibilities, such as cultural facts that one is aware of or has been exposed to.

The parochial sensibility, however, is not sufficient for the application of a concept. The capacity that we gain from having the parochial sensibility that we do cannot be to recognise what is and is not part of a concept's extension full stop, as there is no such thing as the concept's extension full stop.

⁵⁷⁵ A similar sentiment seems to be expressed by Wittgenstein (1953: S217) When considering rule following that there is a certain point at which '...I am inclined to say: "This is simply what I do."' See McDowell (1984a) for discussion.

⁵⁷⁶ Travis 2018: 82

Rather, they can only fix an extension on an occasion, fixing an extension which is most reasonably taken to be useful for that purpose.

The perspective that we have is one that grants us the capacity to recognise when states of affairs should make a given thought true or not, whilst considering the purpose for which we are entertaining that thought. It is a capacity that we might share with others and this allows us to agree with each other about these judgements.

To conclude this section, a judgement is a necessary condition on a thought's being determinate. For these judgements to be made correctly, one needs to judge as a reasonable judge would. These judgements determine a thought's extension on an occasion. Reasonable judgement itself has, at least, two necessary conditions. There must be a purpose in making that judgement. There must also be a parochial sensibility that allows us to connect our thoughts to things in the world. The purpose directs us towards the relevant "openings" of a concept and ensures that there are consequences to judging mistakenly. The parochial sensibility is a capacity to recognise when things in the world are as a given thought describes them. These are two of the necessary conditions on reasonable judgements. Others that I have mentioned include that the judgement must be in pursuit of what is true, such that one feels forced to conclude that the world is that way. One also needs to be familiar with some previous, correct, applications of the concept that one is applying. In the next section I will consider how this view of judgements as extension determining may provide a response to the regress argument. I will then consider some objections to this view.

5.4 The Regress Argument

Fodor and other proponents of meaning determinism, might still be able to use the regress argument to argue against contrary views and (b), the claim that All (or almost all) Sentences in thought are necessarily context-sensitive. In this section I will recap this argument and discuss how Travis's judgement-based view may help to overcome it. The argument goes as follows:

8. Foundational thoughts should not be true in virtue of the truth of any other thought. (If any thought M is true in virtue of the truth of another thought Mn, then M is not foundational).
9. If a thought (tokened sentence in Mentalese) M1 is context-sensitive, then there must be a process by which the context affects the thought.
10. This process of removing context-sensitivity will result in a new thought, M2, which takes the context into account.
11. M1 is true in virtue of the truth of M2 (in that context). (3)
12. All thoughts are context-sensitive (RRC).
13. Therefore, M1 was not foundational. (1, 4)
14. M2 was not foundational (1, 3, 5).

This process can be reapplied to any thought Mn, launching the RRC on a regress of non-foundational thoughts.

Travis's position allows us to deny premise 3. According to Travis's position, a new judgement does not result in a new thought. M1 is true in virtue of how it is judged, rather than its resulting in some new thought. If this is correct, then the regress never begins. In cases where judgement is more demanding, there may be a process of reasoning which results in a thought being tokened. This is, however, just another kind of judgement and that thought will have a determinate truth value in that case. It will not require an additional thought.

Where Fodor and Carston take the bedrock of representation to be thought, which allows for natural language to have the context-sensitivity it does, Travis claims thoughts alone are insufficient for a determinate content. Instead of determinate content coming from thoughts alone he claims that content should stem from the representation's relationship to action. Forging this relationship requires a judgement to be made. These judgements allow for a thought to have its extension be determined on that occasion. This view denies that thoughts are the bedrock of representation. A

thought will not represent in the absence of a judgement. Concepts need applying to objects in the world for them to represent. Once applied they do not need to be reapplied.

The judgement keeps the thought constant but changes the extension that it might have. More than that, acts of judgement are what gives a thought its purchase on the world. Judgements can bring the objects that we encounter in the world into relation to the concepts that we have available to us.

5.4.1 Do Concepts Matter?

On this view there are concepts whose extension might vary with contextual features, such as the purpose which one has for tokening that concept. On Travis's view this comes to the extension being determined by reasonable judges' judgements. This contrasts with Fodor's view according to which the concept was all there was to determining the extension of a concept. Indeed, concepts were, in part, defined by their extension. This method of counting concepts does not seem to be available to Travis as the extension of a given concept can vary. So, we should wonder how to define concepts on Travis's view.

This becomes difficult when we consider two extremes that Travis's view needs to navigate between. On one hand, concepts cannot determine an extension. On the other hand, they cannot do nothing. If concepts did not in some way constrain the judgements that can be made using them then the use of any given concept to express a given content would be acceptable. However, this would make the use of concepts totally arbitrary. Not only is this unpalatable, it is incompatible with a view according to which we can become familiar with certain concepts and not others so that we are able to be reasonable judges for some but not all of them. If the concepts did not matter, then it is hard to see how we could make sense of this. The concepts need to be open textured (in Travis's sense) without being completely open. The danger is that in avoiding the regress we are left with a view according to which the concepts that we use do not matter. If that is so, then context-sensitivity has gone too far.

To put the problem another way, we might decide to understand concepts in a similar way to language, so that concepts provide a character and the character of a thought is a function of its

parts and the way they are combined.⁵⁷⁷ These characters then determine a given content based on the context of use. That content is, on the RRC view considered here, a partial function. The token-underdeterminacy is extreme. So, any new inputs are not going to give an output for that function without a judgement. We make use of the purpose for employing that concept in addition to a human perspective to determine the output of that content on that occasion. But here the content must still do something, as can be seen in cases where the perspective, purpose and possible world are fixed with different outputs for different contents. Compare THERE IS MEAT ON THE CARPET with THERE ARE VEGETABLES ON THE CARPET. If meat is the only thing on the carpet, one is human and is wondering whether to clean the carpet, there should still be a difference in truth-value for these thoughts. So, it seems that the content must play a role.

So, the content needs to be understood in such a way that it can perform a role in producing an output whilst not playing so much of a role that it has a fixed extension. If it cannot perform a role in determining the output, then the content does not seem to play any role in representation. If that is true, then the view faces numerous problems. As suggested above, perspective and purpose are not sufficient for determining an extension. Furthermore, if the content does not play any role then concepts do not seem to be properly distinguishable from one another. Whilst they may have different characters, which we can distinguish intensionally, the resultant contents would not be so distinguishable. So, the use of any one concept would be as good as any other in terms of what they can be used to represent (the concepts would not be *semantically* distinct from one another).

On the other hand, if the contents do determine extensions then they are not token-underdetermined enough. If the contents can produce an output without the assistance of a perspective and purpose, then it is not token-underdetermined. In which case this understanding of

⁵⁷⁷ Maintaining a compositionality of meaning/character whilst denying that there is a compositionality of content or truth conditions is a standard move for radical and Really Radical Contextualists as it allows for an account of what is systematic and productive in language/thought without committing to the idea that the context plays no extensive role in determining the truth-conditions of utterances or tokenings of thought. See Carston 2002: 73-74, Jaque 2017: 83

(b) is not true, and we are not able to use that content to show that underdetermined thoughts can be made determinate (on an occasion) without relying on some other determinate thought. In which case the regress looms large.

Now the issue seems to be how to make sense of a content that does not determine an extension for future inputs which is still able to play a significant role in determining an extension. This is particularly pressing as without an idea of what this role is, and how acts of judging are to interact with it, then it may seem that either the view avoids the regress at the cost of being plausible or it avoids the regress by no longer holding the view that there is underdeterminacy. Neither option is appealing for defenders of RRC.⁵⁷⁸

We can respond as follows. There only needs to be enough of a function to a content that one can be considered reasonable or not to take it to have a given extension on an occasion. The idea would be that in learning our concepts we were exposed to various inputs and learned which output is the correct one for that input. For instance, we were shown cats and dogs and taught that the former can make IS A CAT true whilst the latter makes it false. When it comes to future cases, we make a judgement as to whether the input we currently have would make that thought true or false. So, we have an ability to distinguish concepts based on what we already know to have been their output.

We can make use of previous experience that we have had in learning or acquiring concepts to try to distinguish the contents of those concepts.

⁵⁷⁸ From an exegetical point of view, Travis would likely not accept that this is a problem we should be trying to answer. For him, propositions are primitive, and thoughts will only have content in virtue of their relationship to these propositions. Travis (2018:59) puts the view as follows, were “thought” can be read as “proposition” rather than Mentalese sentence ‘For a thought there is no one set of units out of which it is ultimately built; nor, a fortiori, any one structure which is the way it is (ultimately) constructed. Rather, what is an element of a thought on one way of representing (or presenting) it may not be on another’. Travis (2018: 72) has also claimed that a thought in Frege’s sense is always what is liable to be decomposable in many alternative ways. Each of these makes thought-elements arise. No one of these can claim priority. Thoughts, so thought of, fail to fit the Lego model: there is no set of units which are the ones from which it is built’. Together, these views on thoughts suggest that there is not going to be much that is inherent to the concept that determines its identity. Rather, concepts are related to a proposition and that proposition is occasion sensitive. I will not make use of this view because it is difficult to understand how to make sense of the productivity and systematicity of thought, see Fodor (2005: 87 – 88) for details of these concepts.

This is like Meaning Eliminativism, a view on which we do not make use of an intermediary meaning, such as content or character.⁵⁷⁹ Instead one makes use of the previous uses of an expression and makes a judgement about what the extension of that expression should be on that occasion. On this view, it is the previous uses of that concept that provides a means of distinguishing that concept from other concepts.

This response is in keeping with Travis's reasonable judge line and gives us some criteria by which to distinguish concepts. However, this response seems to fall victim to Kripke-style rule-following considerations. A given extension may be compatible with a range of other alternative concepts, as in the case of "+". Kripke presents the case as follows: '...perhaps in the past I used "plus" and "+" to denote a function which I will call "quus" and symbolize by " \oplus ". It is defined by:

$$x \oplus y = x + y, \text{ if } x, y < 57$$

$$= 5 \quad \text{otherwise.}$$

Who is to say that this is not the function I previously meant by "+"?⁵⁸⁰ Here, the function quus provides the same results as addition would, up until adding numbers equal to or greater than 57. In those cases, the result will only ever be 5. We will suppose that we have never added numbers greater than 57 together before. The problem is that both quus and plus are compatible with past uses of this symbol. When we encounter this new mathematical territory, it is not clear what makes the symbol "+" plus rather than quus.

Non-reductive, "flat-footed", responses are often made to this view. Wright, for example, wonders why 'It is not acceptable, apparently, if the interlocutor claims to recall precisely that [they meant plus by "+", not quus]'.⁵⁸¹ If this response is acceptable, then a concept would still have a determinate content. However, a flat-footed response will not be of great help here as we want to

⁵⁷⁹ Recanati 2004: 146 – 151

⁵⁸⁰ Kripke 1989: 8-9

⁵⁸¹ Wright 2001: 176

know to what we are related to and not what fact about us puts us into that relation.⁵⁸² Knowing that I meant *Z* by a given mentalese symbol “*x*” does not help to distinguish *Z* from *Y*, when *Z* and *Y* seem to be otherwise indistinct. It may even appear that you may as well have meant *Y* for all the difference it makes, and this is the problem now encountered. Concepts still appear arbitrary here.

A response might require us to reject the need for an account of concepts that allows for token-underdeterminacy without falling into Kripke style worries. It seems that expecting contents to semantically distinguish themselves is exactly what one should not expect on this kind of a Travis-like view. That is the role of judgements. This moves the role of distinguishing concepts onto the characters of those concepts (character is a function from contexts of use to contents). For example, a character *C* might, for a context of use *U*, give a content *A*. *A* is then judged true or false for a given input. Suppose it is judged true. This helps to distinguish *A* from *B*, where *B* would have been false for that world. The difference between *A* and *B* may reside in how they have been judged to apply previously (either by oneself or by other reasonable judgers of *As* and *Bs*) and in the characters that are used to get to them. On the judgement dependent view, it is a mistake to think that there can be anything else that distinguishes them. Their content alone will not determine the future extensions. Yet thoughts and concepts can be individuated by their character.

To elaborate, the character of a concept is a function from contexts of use to contents. Each concept has a character that provides the concept with a content on an occasion of use. These characters can be used to help distinguish concepts. For instance, the concept *CAT* will have a character that will make the output the content *cat*. This makes it unlike the concept *DOG*, whose character will give the content *dog* as its output. When we have judged that the content of one is distinct from the content of the other on one occasion, e.g. judging that *x* is a cat and *x* is not a dog, that gives us some means

⁵⁸² Wright himself cashes out his understanding in terms of our ability to make extension determining judgements under certain C-conditions, so that if these C-conditions are met one cannot be wrong about making these judgements (Wright 2001: 192-196). The claim is that linguistic judgements are of this sort (Wright 2001: 212). However, one of Wright’s C-conditions is that one be competent with the concept (Wright 2001: 193). The question here is to do with what the concept is.

of distinguishing between the two concepts. I am assuming that the circumstances of evaluation are fixed in both cases. This helps to distinguish the two concepts and provide a means of individuating concepts.

Being able to distinguish concepts based on their character is useful as it allows us to build up an understanding of which properties that concept has been taken to apply to and why that concept was used to pick out that property on that occasion. This allows an understanding of the concept to be acquired and how to apply it properly, though it does not guarantee that it will be correctly applied. Knowing that the content *cat* has been applied to a given cat for the purposes of deciding whether it will make a good pet does not help to settle whether the Latin speaking cat counts as a cat. It is this store of information that reasonable judges can use to help them to determine the extension of a content in novel cases. It also provides a means of restraining the judgement to some extent.

The human parochial sensibility also provides some constraints on the uses of concepts, especially given an awareness of how those concepts have been applied in the past. Some applications of a concept will likely just seem wrong to us, for instance, thinking CAESAR A PRIME NUMBER will strike most of us as false (if it does not strike one as nonsense). The parochial sensibility and awareness of past uses do not serve to make the content determinate, but they do help to constrain which concepts reasonable judges will take to be correctly applied to a given situation. The actual extension will still be determined by judgement. If Travis's position is correct, we cannot expect the content of a concept to determine its extension independently of these factors.

These judgements are not arbitrary because there can still be cases in which one can judge incorrectly.⁵⁸³ This is in those cases in which one's judgements are not in accord with those of a reasonable judge. So, if one judges that a dog is a cat one is not acting as a reasonable judge would.

⁵⁸³ Wright (2001: 194) takes a similar approach. If a judgement does not determine the correct extension, then that judgement cannot have been made under C-conditions.

One has not met the criteria for making an extension determining judgement. This is only an attempt at judging, and it has failed. The selection of the concept DOG, instead of CAT, would have bought one's judgement in line with that of a reasonable judge. This failure to judge correctly may be due to a lack of familiarity with the concepts or failing to meet one of Travis's other C-conditions.

This approach has some advantages. For one, it preserves compositionality of character and thereby allows for distinguishing concepts from one another. Each concept has a distinctive character, and this may help to distinguish concepts. Though it does not help to distinguish them in virtue of their contents. It also rejects the need for content to perform the role of determining extension, with that role instead falling to judgement.

This is not, however, without some cost. For one thing, the resulting picture is much less neat than Fodor's. Rather than having only a circumstance of evaluation and a content, we now also have a human perspective and a purpose to consider in making a judgement. It is not obvious how best to understand these parts. The human perspective might seem to be too messy an idea to be useful. This is not a flaw unique to Travis's position, as Fodor seems to rely on something similar. Whilst considering these views in opposition to one another, this point doesn't seem to favour one over the other. Fodor's overall picture is neater but cannot account for changes in extension.

5.5 Conclusion

In this chapter I have presented an understanding of Travis's RRC which focused on token-underdeterminacy. Whilst this is an important strand of Travis's work, I do not take this to be Travis's actual position, rather it is a Travis-like view. To make this view clearer, I have contrasted the view with Fodor's. In section 5.1, I presented a reading of Fodor according to which the extension of a concept is determined entirely by the identity of that concept. On Fodor's view mental representations are the primitive bearers of intentional content. On this view, the mental state determines the extension of a concept.

In section 5.2 I presented one criticism of Fodor's position. This was the rule-following argument which aims to undermine the idea that a given mental state can determine its extension in all future cases. Whilst Travis's arguments seem to undermine Fodor's position, as Fodor's view of categories as including all actual and possible extensions of a concept seems to have implausible consequences. They also suggest that a content alone does not determine whether that concept can be applied correctly in a given situation. For avoiding a regress, it is sufficient that there be some concepts that are radically underdetermined and that these thoughts are not replaced by further underdetermined thoughts.

In section 5.3 I sketched the Travis view. On this view concepts are token-underdetermined, so they do not determine their extension. Furthermore, thoughts are not the primitive bearers of intentional content on Travis's view. Instead, it takes acts of judgement to make a Mentalese sentence about anything in the world. These judgements are occasion specific and do not resolve all the token-underdeterminacy that that concept has. Rather, they are for applying that concept on that occasion.

On the Travis view, a thinker must be competent with a concept and there must be a purpose for applying that concept. There must also be a parochial sensibility from which the judger makes their judgement. The purpose is necessary for directing us to the open texture that is relevant for that case. It also helps us to decide whether this concept is useful for what we have in mind. The perspective grants us capacities to recognise when something is or is not the case. Finally, a judgement must be made either by a reasonable judge or be in line with those of a reasonable judger. These judgements can be more or less involving on the part of the agent.

In section 5.4, I applied this view of judgements to the regress argument as suggested by Carston and Fodor. On this argument, underdeterminacy is unavoidable for RRC. On the judgement dependent view, that a thought is token-underdetermined is not an issue. Thoughts are not the primitive bearers of intentional content, judgements are. These judgements make the thought true or false for that input. Whilst the judgement does not determine all future outputs for the range of possible

inputs for a given thought, it does enough for those cases in which judgement is successful. So, whilst we may have to make a lot of judgements, it does not seem that the underdeterminacy left over after an act of judging is a problem. The concept gets applied and requires no further action. So, this view avoids a regress.

To finish section 5.4, I considered whether such a view was still plausible. To do so I considered whether there was a threat to the identity of concepts. If concepts were all open textured in the radical way that Travis endorses, then it seems that any concept could be used in the place of any other. In which case, there is an unappealing arbitrariness to our concepts that we should find troubling. In response to this I suggested that we identify concepts by their character and not seek any properly semantic identification of concepts. That is precisely what we should not expect them to do.

The defence of this Travis view provides a third way in which thoughts might be context-sensitive. This context-sensitivity is due to token-underdeterminacy. This view presents the most radical version of context-sensitivity of thought that I have considered in the thesis. It requires a view of truth as judgement dependent as well as requiring token-underdeterminacy of content. It does, however, avoid the regress argument and thereby provide evidence for the claim that thoughts can be context-sensitive.

6. Conclusion

The RRC view states that all (or almost all) thoughts are context-sensitive. The most serious challenge to this view is the regress argument, according to which any context-sensitivity in thought is purely contingent and must evaporate on analysis. The argument goes as follows:

1. Foundational thoughts are not true in virtue of the truth of any other thought. (If any thought M is true in virtue of the truth of another thought Mn, then M is not foundational).
2. If a thought (=tokened sentence in Mentalese) M1 is context-sensitive, then there must be a process by which the context affects the thought.
3. This process of removing context-sensitivity will result in a new thought, M2, which takes the context into account.
4. M1 is true in virtue of the truth of M2 (in that context).
5. All thoughts are context-sensitive (RRC).
6. Therefore, M1 was not foundational. (1, 4)
7. M2 was not foundational (1, 3, 5).

This process can go on infinitely, with any thought Mn failing to be foundational (a thought is foundational when its truth does not depend on the truth of any other). That would be a serious problem for RRC as the underdeterminacy involved would not be removed. The responses that I have presented have all denied that premise 3 is true. This means that each of the responses to this argument that I have considered have attempted to show how an underdetermined thought can be sensitive to the context without requiring that an additional thought be created.

As the regress argument was an argument intended to show that context-sensitivity at the level of foundational thought is impossible, responses to this argument show that it is possible for a foundational thought to be context-sensitive. I take this to be a good reason to endorse *weak*

foundational context-sensitivity (modal), the claim that some foundational sentences in thought *could* be context-sensitive.

One of the key motivations for endorsing the Mixed View (MV), according to which all (or almost all) natural language sentences are context-sensitive (claim (a)) whilst denying thoughts are context-sensitive, was that it seemed to be impossible for thoughts to be context-sensitive. So that *weak foundational context-sensitivity (modal)* is true undermines this key motive for endorsing the MV over the RRC view.

I do not, however, take myself to have argued that the MV is false, or that all thoughts are, or must be, context-sensitive. An exception to this is the case of *de se attitudes*, which seem to require indexicals in thought. The default position seems to be that if all or almost all natural language sentences are true then the same will be true of thoughts. If one accepts this assumption, then the counter argument that I have presented in favour of the possibility of context-sensitive thoughts will help to convince those who accept (a) that (b) is also true. Thus, this is an assumption that proponents of the Mixed View should be keen to deny. They might deny that there are reasons to suppose that we can infer the nature of thoughts from the nature of language. This would be to deny that the two are analogous in these respects. I have only briefly touched on this issue. The prospects for establishing an analogy between thoughts and natural language are unclear.

It is worth noting, however, that comparing thoughts to natural language sentences is a useful tool and one that proponents of the MV have made use of. For instance, endorsing the language of thought hypothesis (LOTH) has been useful for explaining how thought might be systematic and productive. It is unlikely that they would want to deny that there is any analogy to be drawn here.

I have also not argued that (a) is true, I have instead assumed that it is true as a point of common ground between the MV and RRC. However, if (a) were false then it would be likely that (b) is also false. This does not mean that (a) is not controversial and there is an ongoing debate between Minimal Semanticists who deny that it is true and Radical Contextualists who affirm it.

These are some of the obstacles to proving that the RRC view is correct and I do not consider myself to have removed these obstacles. Whilst I have presented some arguments in favour of RRC, many of these rely on comparing thoughts to natural language sentences. If it could be shown that thoughts are disanalogous to natural language sentences, then these arguments may not hold up so well.

Nevertheless, the argument that thoughts can be context-sensitive means that these arguments should get more attention. This possibility can no longer be dismissed. It also opens the possibility that thoughts are not the bedrock of human representation. Instead, it is possible that it is thoughts in conjunction with some other factor that enables thoughts to express propositions. It cannot just be assumed that thoughts will be able to represent independently of the role that they play in our lives.

I have also presented arguments in favour of the context-sensitivity of thought being compatible with the computational theory of mind (CTM) in Chapter 2. I have made this argument on the basis that a CTM cannot be thoroughly computational unless the Frame Problem has been solved. The Frame Problem suggests that a lot of reasoning cannot be done computationally, as computations cannot sort information into sets of relevant information. However, the Frame Problem allows for the possibility that there are still parts of the mind that are computational, provided that this computation occurs within a mental module.

Context-sensitivity seems incompatible with computation as it violates the type sensitivity constraint (the TSC, the claim that if two representations differ in their content then they must be of different types) by treating different tokens of a type in different ways. Therefore, context-sensitivity cannot be treated computationally. However, this does not make it incompatible with the CTM. It just means that these processes are not computational and occur in a general, non-computational, reasoning system.

Whilst this is a steep price to pay, I believe that it is one that we are already paying because of the Frame Problem. Fodor, one of the key proponents of the CTM, is happy to endorse this kind of

cognitive architecture. Furthermore, this compatibility with the CTM is important, even when only some of the mind can be computational, the CTM is one of the leading accounts in cognitive science. The Frame Problem also highlights that relevance is a difficult property to understand within a CTM, though the CTM has the advantage that some cognitive processes can be understood. That relevance is difficult to understand has put the MV and the RRC view on more equal footing as it is something that they both struggle to account for.

In Chapter Three I discussed the possibility of *de se attitudes* being indexical thoughts that are necessarily context-sensitive. Here I recounted Ninan's arguments that *de se attitudes* do present their own unique problem and are not merely a variant on Frege cases of informative identity statements. I then considered an argument given by Ball (though not attributed to him) to the effect that indexicals are not possible on a CTM. I then argued that *de se attitudes* are best accounted for as indexicals. Other models do not make it apparent that the referent is oneself.

As *de se attitudes* seem to be necessarily indexical and cannot be replaced by non-*de se attitudes* whilst maintaining the same cognitive effects, it follows that there are going to be indexical thoughts. So, these thoughts are both context-sensitive and foundational. In which case, there can be some foundational context thoughts. This was an instance of what I have called character-underdeterminacy. A representation is character-underdetermined only if a component of that representation makes an explicit reference to context to establish content.

In Chapter Four I considered Searle's Background theory as a version of type-underdeterminacy at the level of thought. On this view, we make use of a range of Background assumptions (that need not themselves be representations) to put type-underdetermined representations into a relation with a content. Whilst this presents a possible RRC view, it does not yet do enough to resist the regress argument as it does not say how the Background should bring about determinacy. If it were by generating another thought, then there would be a regress. To respond to this concern, I suggested

that Perry's concerning relation is a useful way of understanding how the Background can make type-underdetermined thoughts determinate.

The concerning relation is a relationship that exists between a representation and that representation's content. Crucially, the concerning relationship does not require an additional thought or any changes to the representation itself. Instead, the Background allows a representation to play a role in a person's life. The suggestion then is that the Background can bring representations into a concerning relation with a particular content. This means that the thought can be made determinate without requiring an additional thought to be tokened. Hence this view also avoids a regress.

This view encounters some problems that might affect its plausibility. I discussed the objection that this view leads to holism which would make the possibility of agreeing and disagreeing with others and oneself across time impossible. I argued that this objection places the bar on agreement and disagreement too high. If we accept a lower threshold for agreement and disagreement, then there is no concern here. Agreement and disagreement remain possible.

In Chapter Five I considered a version of token-underdeterminacy, according to which content only gives a partial function from possible worlds to truth values. This is a version of Travis's view according to which future instances of objects falling under a concept are not determined by that concept. This is a radical view and, to make it clearer, I presented Fodor's view of concepts to provide a foil to the Travis-like view I develop here. According to Fodor's view, the extension of a concept is determined by actual and possible causal relationships between the extension of a concept and tokenings of that concept. I gave an argument against Fodor's view before considering a Travis-like view.

On the view that I developed, the extension of a concept is dependent on the judgements made by a reasonable judge. These judgements determine the extension of a concept on a case by case basis. These acts of judgement make it the case that a token-underdetermined thought can be made

determinate in that instance. A further thought is not required to make this thought determinate. In which case, it is a view that is also capable of avoiding a regress.

The ability to make these judgements is dependent on several conditions being met. I have not given an exhaustive list but have focused on two crucial ones. Firstly, there needs to be a purpose to that representation being tokened. Secondly, the judge must have a human parochial sensibility. Both features need to be present for a reasonable judgement to be made and they can both influence verdict of the judgement. For instance, if the purpose were to change then so might the judgement as is illustrated in the case of THE LEAVES ARE GREEN.

One concern with this view is that it takes a very radical view truth, according to which a proposition is made true or false depending on the judgement that it is. I have attempted to alleviate this by arguing that judgements are not trivial and can require a lot of additional information and familiarity both with the concept to be applied and with the object concerned. For the reasonable judge, the judgements they make should not be arbitrary or a matter of opinion. Instead the verdict that they reach should seem to be forced upon them by the circumstances. This should go some of the way towards showing that whilst truth is judgement-dependent, it is not an arbitrary or simple matter. It may also be unclear how to decide who is and is not a reasonable judge in a given case.

Another concern with the Travis-like view is it makes concepts too thin to be properly distinguishable from each other. In response to this concern I have argued that we should not expect the content to distinguish concepts on this view, instead we need to rely on character to do so. Once concepts are distinguished in this way, we can associate different bodies of information with them, including what they have been applied to in the past and why, to make the selection of concepts non-arbitrary. Furthermore, these concepts can still be applied incorrectly. This suggests that a token-underdetermined account of thoughts is also possible.

The method that I have pursued in this thesis has been to consider a view that exhibits one of the kinds of underdeterminacy that I am interested in and whether they can avoid a regress. The focus

has been on whether they can take the context into account without generating a further thought. This is not meant to be an exhaustive list of the possibilities or to say that these are the most plausible alternatives. However, they are the ones that one sees most often in the literature, without consideration of how they might deal with this regress argument. Alternative versions of RRC will also have to take this approach of denying premise 3, whilst avoiding other problems and it is likely that they will encounter similar problems to those that I have discussed in each case.

A drawback of this method is that it is not easy to consider which of these versions of underdeterminacy is the most appealing, or how they will relate to one another. In terms of how they relate to one another, there are several possibilities. It may be that they are all compatible with one another, though that would require a lot of additional cognitive work which would make that unappealing. Each version makes use of different mechanisms which won't always be making use of the same contextual relation. Such a view would require a thought be put into a concerning relation with a given content and then that content be judged to resolve token-underdeterminacy. This would be taxing and complex. So, a future project would be to consider which of the Background view and Travis-like view is the strongest.

In summary, I believe that I have made the following contributions in this thesis. The most general contribution is that I have given a detailed account of how RRC views might deny premise 3 of the regress argument, of which I have presented my own version. I have offered a defence of the Frame Problem from counter arguments presented by Chow and Samuels in Chapter Two. This was used to motivate a view of cognitive architecture according to which modules are computational, but a general processing system is not. I have also presented a defence of mental modules from numerous objections. This is important for supporting a theory of mind that allows for context-sensitive thoughts.

I have also given a counter argument to Argument 1 that indexicality is incompatible with the CTM in chapter 3. This response to Argument 1 has relied heavily on the Frame Problem and a modular

account of the mind. In conjunction with *de se indexicalism*, this suggests that there are context-sensitive thoughts.

In Chapter Four I have discussed how Searle's view might be conjoined with Perry's view of a representation concerning different states of affairs via being in a relation with different contents. I have also discussed some of the responses that might be made to this view. This has the advantage of spelling out some of these issues with a concerning account and responding to them. The concerning relation is sometimes taken for granted in the literature and this helps to make the case for it more persuasive.

In Chapter Five I have attempted to show how Fodor and Travis's views contrast with one another, as well as some points at which they seem to hold similar views. I have also developed my own version of Travis's view, as one on which the extension of a concept is dependent on reasonable judges' judgements. This shows how radical RRC can be, as truth is dependent on acts of judgement on this view. These views suggest that RRC is a viable alternative to the MV and removes one of the key objections to the RRC as made by proponents of the MV. Therefore, it is possible that thoughts are context-sensitive.

I've shown that, contrary to a common view, the regress argument is not a knock down argument against context-sensitive thought. I've argued that what it shows rather that any understanding of context-sensitive thought must avoid the idea that context-sensitivity introduces another thought. In response to this realisation, I have shown that this can be avoided. Furthermore, it can be avoided whether we take context-sensitivity to take the form of character-underdeterminacy, type-underdeterminacy or token-underdeterminacy. This matters because it changes how we might understand representation and how to understand our own cognition. In place of atomic representations, we can make use of representations that depend on the context in some way to play their role.

Bibliography

- Allott (2010) *Key Terms in Pragmatics* London: Continuum International Publishing Group
- Allott and Textor (2012) "Lexical Pragmatic Adjustment and the Nature of *Ad hoc* Concepts" in *International Review of Pragmatics* 4 pp. 185-208
- Alshwari H. (1987) *Memory and Context for Language Interpretation* Cambridge: The Press Syndicate for the University of Cambridge
- Anscombe GEM (1975) "The First Person" in (ed.) Guttenplan *The Philosophy of Mind: Wolfson College Lectures 1974* pp. 22 – 36
- Apperly et al. (2009) "Do Humans Have Two Systems to Track Beliefs and Belief-Like States?" in *Psychological Review* Vol. 116, No. 4 pp. 953 – 970
- Atlas (2007) "Meanings, Propositions, Context, and Semantical Underdeterminacy" in (ed) Gerhard *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics* Oxford: Oxford University Press
- Austin J. (1975) *How to do Things with Words* Oxford: Clarendon Press
- Adydede (1999) "On the Type/Token Relation of Mental Representations" in *Facta Philosophica: International Journal for Contemporary Philosophy* v.7/12/1999
- Aydede (2010) "The language of Thought Hypothesis" *The Stanford Encyclopaedia of Philosophy* (ed.) E. Zalta
- Balcetis and Dunning (2006) "See What You Want to See: Motivational Influences on Visual Perception" in *Journal of Personality and Social Psychology* 91 pp. 612-625
- Ball (2010) "Can There Be Indexical Thoughts" Draft www.st-andrews.ac.uk/~db71/ball_indexical_thoughts.pdf
- Ball (2015) "Indexicality, Transparency, and Mental Files" in *Inquiry* Vol. 58. No 4 pp. 353 - 367
- Barret and Kurzban (2006) "Modularity in Cognition: Framing the Debate" in *Psychological Review* Vol. 113, No. 3 pp. 628 – 647
- Bermudez (2017) *Understanding I: Language and Thought* Oxford: Oxford University Press
- Bilgrami (1998) "Why Holism is Harmless and Necessary" in *Language, Mind and Ontology* Vol. 12 pp. 105 – 126
- Billon and Guillot (2014) "Can Fregeans have "I"-Thoughts?" in *Revista de Filosofía de la Universidad de Costa Rica* Vol. 136 pp. 97 – 105
- Borg (2000) "Complex Demonstratives" in *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* Vol. 97. No. 2 pp. 229 – 249.
- Borg (2001) "An Expedition Abroad: Metaphor, Thought, and Reporting" in *Midwest Studies in Philosophy* Vol. XXV pp. 227 - 248
- Borg (2001a) "The Metaphysics and Epistemology of Singular Terms" in *Philosophical Papers* Vol. 30:1 pp.1 – 30

- Borg (2004) *Minimal Semantics* Oxford: Oxford University Press
- Borg (2009) "Must a Semantic Minimalist Be a Semantic Internalist?" in *Proceedings of the Aristotelian Society, Supplementary Volumes*, Vol. 83 (2009), pp. 31 – 51
- Borg (2012) *Pursuing Meaning* Oxford: Oxford University Press
- Borg (2014) "From Ape Empathy to Human Morality?" in *Analysis* Vol.74 No. 4 pp. 577 – 587
- Borg (2016) "Exploding Explicatures" in *Mind and Language* Vol. 31 No. 3 pp.335 – 355.
- Borg (2017) "Explanatory Roles for Minimal Content" in *Nous* ... pp. 1 – 27
- Borg (2018) "On Deflationary Accounts of Human Action Understanding" in *Review of Philosophy and Psychology* Vol. 9, Issue 3, pp. 503 – 522
- Braun, David, "Indexicals", *The Stanford Encyclopaedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2017/entries/indexicals/>>.
- Brown, Curtis, "Narrow Mental Content", *The Stanford Encyclopaedia of Philosophy* (Summer 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2016/entries/content-narrow/>>.
- Cain (2016) *The Philosophy of Cognitive Science* Cambridge: Polity Press
- Cappelen (2008) "The Creative Interpreter: Content Relativism and Assertion" in *Philosophical Perspectives* Vol. 22, Philosophy of Language (2008) pp. 23 – 46
- Cappelen and Dever (2013) *The Inessential Indexical: On the Philosophical Insignificance of Perspective and the First Person* Oxford: Oxford University Press
- Cappelen and Lepore (2005) *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism* Oxford: Blackwell Publishing Ltd
- Carl (1994) *Frege's Theory of Sense and Reference: Its Origins and Scope* Cambridge: Cambridge University Press
- Carruthers (2002b) "The cognitive functions of language" in *Behavioural and Brain Sciences* Vol. 25, pp. 657-726
- Carruthers (2006) *The Architecture of the Mind* Oxford: Oxford University Press
- Carruthers (2009) "Language in Cognition" in (eds) E. Margolis, R. Samuels and S. Stich *The Oxford Handbook of Philosophy of Cognitive Science* Oxford: Oxford University Press
- Carston (2002) *Thoughts and Utterances: The Pragmatics of Explicit Communication*, Malden: Blackwell Publishers
- Carston and Hall (2012) "Implicature and Explicature" from ResearchGate at https://www.researchgate.net/publication/265483734_Implicature_and_Explicature also in (ed.) Hans-Jörg Schmid *cognitive Pragmatics* pp. 47 – 84
- Carston and Powell (2008) "Relevance Theory: New Directions and Developments" in (ed.) E. Lepore and B. Smith *The Oxford Handbook of Philosophy of Language* Oxford: Oxford University Press
- Catañeda (1968) "On the Logic of Attributions of Self-Knowledge to Others" in *The Journal of Philosophy* Vol. 65 No. 15 pp. 439 – 456

- Chalmers (2018) "Juhani Yli-Vakkuri and John Hawthorne, *Narrow Content*, Oxford University Press" in <https://ndpr.nd.edu/news/narrow-content/>
- Chomsky, N. (1986) *Knowledge of Language: Its Nature, Origin, and Use* Westport Connecticut, Praeger Publishers
- Chomsky (2006) *Language and Mind* (3rd edition) Cambridge: Cambridge University Press
- Chow (2013) "What's the Problem with the Frame Problem?" in *Review of Philosophy and Psychology* Vol. 4 pp. 300 – 331
- Church (1941) *The Calculi of Lambda-conversion* Princeton: Princeton University Press
- Collins (2017) "Faculties and Modules: Chomsky on Cognitive Architecture," in McGilvray, J. (ed.) *The Cambridge Companion to Chomsky*. 2nd edn. Cambridge: Cambridge University Press, pp. 217–234. doi: 10.1017/9781316716694.011.
- Collins and Dobler (2018) "The Work of Charles Travis: A Contemporary Overview" in *The Philosophy of Charles Travis: Language, Thought and Perception* Oxford: Oxford University Press
- Colyvan, Mark, "Indispensability Arguments in the Philosophy of Mathematics", *The Stanford Encyclopedia of Philosophy* (Spring 2015 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2015/entries/mathphil-indis/>>.
- Corazza and Whitsey (2003) "Indexicals, Fictions, and Ficta" in *Dialectica* Vol. 57, No. 2 pp. 121 – 136
- Cosmides and Tooby (1994) "Origins of Domain Specificity: The Evolution of Functional Organization" in (eds.) Hirschfeld and Gelman *Mapping the Mind: Domain Specificity in Cognition and Culture* Cambridge: Cambridge University Press
- Cosmides, Barrett and Tooby (2010) "Adaptive Specializations, Social Exchange, and the Evolution of Human Intelligence" in *Proceedings of the National Academy of Sciences of the United States of America* Vol. 107 no. 2 pp. 9007 – 9014
- Crane (2002) "Intentional Objects" in (ed.) Borg *Meaning and Representation* Oxford: Blackwell Publishers
- Cumming, Sam, "Names", *The Stanford Encyclopedia of Philosophy* (Fall 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2016/entries/names/>>.
- Davidson (1967) "Truth in Meaning" in *Synthese: Language in Use Including Wittgenstein's Comments on Frazer and a Symposium on Mood and Language-Games* Vol. 17, No.3 pp. 304 – 323
- Del Campo et al (2016) "Decision Making Styles and the Use of Heuristics in Decision Making" in *Journal of Business Economics* Vol. 86 pp. 389 – 412
- Dennett (1979) *Brainstorms: Philosophical Essays on Mind and Psychology* Hassocks: The Harvester Press Ltd.
- DeRose (1992) "Contextualism and Knowledge Attributions" in *Philosophy and Phenomenological Research* Vol. 52, No. 4 pp. 913 – 929
- De Gaynesford (2006) *I: The Meaning of the First Person Term* Oxford: Oxford University Press
- Dickie (2010) "The Generality of Particular Thought" in *The Philosophical Quarterly* Vol. 60 No. 240 pp. 508-31

- Dobler (2017) "The Occasion-Sensitivity of Thought" available open access at <https://link.springer.com/article/10.1007/s11245-017-9532-2>
- Dummett (1996) *The Seas of Language* Oxford Scholarship Online
- Egan A. (2006) "Secondary Qualities and Self-Location" in *Philosophy and Phenomenological Research* in Vol. LXXII, No. 1
- Egan A. (2007) "Epistemic Modals, Relativism and Assertion" in *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* Vol. 133, No. 1
- Egan F. (2014) "How to Think About Mental Content" in *Philosophical Studies* Vol. 170 pp. 115-135
- Egan F. (2017) "Function-Theoretic Explanation and the Search for Neural Mechanisms" in (ed.) D. Kaplan *Explanation and Integration in Mind and Brain Science* Oxford: Oxford University Press
- Elguardo (2005) "Fodor's Inexplicitness Argument" in (eds.) Werning, Edouard, Schurz *The Compositionality of Meaning and Content. Volume 1: Foundational Issues* Berlin: De Gruyter
- Evans (1981) "Understanding Demonstratives" in (eds.) Parret, Bouveresse *Meaning and Understanding* Berlin: Walter de Gruyter
- Evans (1982) *The Varieties of Reference* Oxford: Oxford University Press
- Ferreira et al. (2009) "Good Enough Language Processing: A Satisficing Approach" in N. A. Taatgen & H. van Rijn (eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society*. pp. 413--418
- Firestone and Scholl (2016) "Cognition Does Not Affect Perception: Evaluation the Evidence for "top-down" effects" in *Behavioural and Brain Sciences* Vol. 39, E229
- Fitch, Greg and Nelson, Michael, "Singular Propositions", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/win2016/entries/propositions-singular/>](https://plato.stanford.edu/archives/win2016/entries/propositions-singular/).
- Fodor (1975) *The Language of Thought* Harvard: Harvard University Press
- Fodor (1983) *Modularity of Mind: An Essay on Faculty Psychology* Massachusetts: MIT Press
- Fodor, J. (1989). "Review Essay: Remnants of Meaning by Stephen Schiffer." in *Philosophy and Phenomenological Research* Vol. 50(2): 409-423
- Fodor (1994) "Concepts: A Potboiler" in *Cognition* Vol. 50 pp. 95 – 113
- Fodor (1998) *Concepts: Where Cognitive Science Went Wrong* Oxford: Clarendon Press
- Fodor (1999) "Diary" in *London Review of Books* Vol. 21 No. 19 pp. 68 – 69 accessed from <https://www.lrb.co.uk/v21/n19/jerry-fodor/diary>
- Fodor (2000) *The Mind Doesn't Work that Way* London: The MIT Press
- Fodor (2001) "Language, Thought and Compositionality" in *Mind and Language* Vol. 16 no. 1 pp. 1 - 15
- Fodor (2005) *Hume Variations* Oxford Scholarship Online
- Fodor (2009) "Enough with the Norms Already" in Heike and Leitgeb (eds.) *Reduction: Between the Mind and the Brain* Ontos Verlag

- Fodor (2010) *LOT2: The Language of Thought Revisited* Oxford: Clarendon Press
- Fodor and Pylyshyn (2015) *Minds without Meanings: An Essay on the Contents of Concepts* London: The MIT Press
- Frankenhuis et al. (2007) "Evolutionary Psychology Versus Fodor: Arguments For and Against the Massive Modularity Hypothesis" in *Philosophical Psychology* Vol. 20 No. 6 pp. 687 - 710
- Frege (1918/2010) "The Thought: A Logical Enquiry" in (ed.) Matrinich *The Philosophy of Language* 5th edition, Oxford: Oxford University Press
- Frege (1892/2010) "On Sense and Nominatum" in (ed.) Matrinich *The Philosophy of Language* 5th edition, Oxford: Oxford University Press
- Frege (1948) "Sense and Reference" in *The Philosophical Review* Vol. 57 No. 3 pp.209 – 230.
- Geurts and Rbio-Fernández (2015) "Pragmatics and Processing" in *Ratio (new Series)* Vol. 28pp. 446 - 469
- Grice (2010/1975) "Logic and Conversation" in (ed.) A.P. Martinich *The Philosophy of Language* Oxford: Oxford University Press
- Goldstein et al (2002) "Why and When Do Simple Heuristics Work?" in (eds) Gigerenzer and Selten *Bounded Rationality: The Adaptive Toolbox* Cambridge: MIT Press
- Grossi (2014) "A module is a Module is a Module: Evolution of Modularity in Evolutionary Psychology" in *Dialect Anthropol* Vol. 38, pp. 333 - 351
- Guerts (2019) "Communication as Commitment Sharing: Speech Acts, Implicatures, Common Ground" in *Theoretical Linguistics* Vol. 45, No. 1-2 pp. 1-30
- Guillot (2016) "Thinking of Oneself as the Thinker: The Concept of Self and the Phenomenology of Intellection" in *Philosophical Explorations: An International Journal for the Philosophy of Mind and Action* Vol. 19:2 pp. 138-160
- Hacker (2013) "The Linguistic Turn in Analytic Philosophy" in *The Oxford Handbook of The History of Analytic Philosophy* Oxford: Oxford University Press
- Hanks (2009) "Recent Work on Propositions" in *Philosophy Compass* pp. 469 - 486
- Hansen (unpublished) "Really Radical Contextualism"
- Hansen (2010) *Radical Contextualism: A Dissertation Submitted to the Faculty of the Division of the Humanities in Candidacy for the Degree of Doctor of Philosophy* Chicago: The University of Chicago
- Hansen (forthcoming) "Just What Is It That Makes Travis's Examples So Different, So Appealing" in (eds) J. Collins, A. Davies and T. Dobler *Themes from Charles Travis: On language, Thought and Perception* Oxford: Oxford University Press, accessed from https://www.academia.edu/6516552/Just_What_Is_It_That_Makes_Traviss_Examples_So_Different_So_Appealing
- Hansen (2018) "Just What Is It That Makes Travis's Examples So Different, So Appealing" in (eds) J. Collins, A. Davies and T. Dobler *Themes from Charles Travis: On language, Thought and Perception* Oxford: Oxford University Press
- Harmen (1999) *Reasoning, Meaning, and Mind* Oxford: Oxford University Press

- Harris (2018) "Semantics without Semantic Content", forthcoming in *Mind and Language*. Accessed from <https://semanticsarchive.net/Archive/2IzNzE4Z/DanielHarris-SemanticsWithoutSemanticContent.pdf>
- Heil (1981) "Does Cognitive Psychology Rest on a Mistake?" in *Mind* Vol. XC, 321-342
- Heil (2004) *Philosophy of Mind: A Contemporary Edition* (2nd ed.) Abingdon: Routledge
- Hermer-Vazquez (2002) "Viewing Cognitive Mechanisms in the Context of Biology" in *Behavioural and Brain Sciences* Vol. 25, pp. 689 – 690
- Hershfield (2005) "Rule Following and the Background" in *Linguistics and Philosophy* Vol. 28, No.3 pp. 269 – 280
- Higginbotham (2013) "Speaking (and Some Thinking) of Oneself" in (eds.) Feit and Capone *Attitudes De Se: Linguistics, Epistemology, Metaphysics* Leland Stanford: CSLI Publications
- Hodges, A. (2001) "Turing" in *The Great Philosophers* in (ed.) F. R. R Monk. London, Weidnfeld & Nicolson
- Hookway, Christopher, "Pragmatism", *The Stanford Encyclopedia of Philosophy* (Summer 2016 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/sum2016/entries/pragmatism/>](https://plato.stanford.edu/archives/sum2016/entries/pragmatism/).
- Hötting and Röder (2004) "Hearing Cheats Touch, but Less in Congenitally Blind than in Sighted Individuals" in *Psychological Science* Vol. 15 No. 1 pp. 60 – 64
- Hume (2013/1757) *Of the Standard of Taste* Birmingham: The Birmingham Free Press
- Huestegge, Pieczykolan, Koch (2014) "Talking While Looking: On the Encapsulation of Output System Representations" in *Cognitive Psychology* Vol. 73 pp. 72 – 91
- Huth et al (2016) "Natural Speech Reveals the Semantic Maps that Tile Human Cerebral Cortex" in *Nature* Vol. 532
- Jackman (2007) "Minimalism, Psychological Reality, Meaning, and Use" in (eds.) Preyer and Peter *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics* Oxford: Oxford University Press
- Jaque (2017) "Are Mental Representations Underdeterminacy-Free?" in *Synthese* Vol.194 (7) pp. 1-22
- Jaszczolt (2013) "Contextualism and Minimalism on De Se Belief Ascription" in (eds.) Feit and Capone *Attitudes De Se: Linguistics, Epistemology, Metaphysics* Leland Stanford: CSLI Publications
- Jylkkä (2008) "Why Fodor's Theory of Concepts Fails" in *Minds and Machines* Vol. 19 pp. 25 – 46
- Kamermans and Schmitts (2004) "The History of the Frame Problem" <https://staff.fnwi.uva.nl/b.bredeweg/pdf/BSc/20032004/KamermansSchmits.pdf>
- Kaplan (1989) "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals" in (ed.) Almog J. *Themes from Kaplan* Oxford: Oxford University Press
- Keane (2018) "Contour Interpolation: A Case Study in Modularity of Mind" in *Cognition* Vol. 174 pp. 1 – 18

- Kecskes (2017) "The Interplay of Recipient Design and Salience in Shaping Speaker's Utterance" in (eds.) Ponte and Korta *Reference and Representation in Thought and Language* Oxford: Oxford University Press
- Khoo (2017) "Code Words in Political Discourse" in *Philosophical Topics* Vol. 45, No.2, pp.33 – 64
- King (2003) "Tense, Modality, and Semantic Values" in *Philosophical Perspectives* Vol. 17, Language and Philosophical Linguistics pp. 195 – 245
- Knoll A., Rey G. (2016) "Anthropod Intentionality" for (ed.) Beck J. *Routledge Handbook of Philosophy of Animal Minds* Abingdon: Routledge
- Kölbel (2009) "The Evidence for Relativism" in *Synthese* Vol. 166, No. 2 pp. 375 – 395
- Korta and Perry (2008) "The Pragmatic circle" in *Synthese* Vol. 165, No. 3 pp. 347 - 357
- Korta and Perry (2007) "Radial Minimalism, Moderate Contextualism" in (ed.) Peter *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics* Oxford: Oxford University Press
- Kuczynski (2007) *Conceptual Atomism and the Computational Theory of Mind: A Defense of Content-Internalism and Semantic Externalism* John Benjamins Publishing Company
- Kusch M. (2006) *A Sceptical Guide to Meaning and Rules: Defending Kripke's Wittgenstein* Bucks: Acumen
- Kripke (1989) *Wittgenstein on Rules and Private Language: An Elementary Exposition* Malden: Blackwell Publishing
- Lappin S. (2003) "Semantics" in (ed.) Mitkov R. *The Oxford Handbook of Computational Linguistics* Oxford: Oxford University Press
- Laurence, S. (1996) "A Chomskian Alternative to Convention-Based Semantics." *Mind* 105(418): 269-301
- Lepore and Loewer (1981) "Translational Semantics" in *Synthese* Vol. 48, No. 1 pp. 121 – 133
- Lepore and Stone (2014) *Imagination and Convention: Distinguishing Grammar and Inference in Language* Oxford: Oxford University Press
- Lewis (1970) "General Semantics" in *Synthese* Vol. 22, No. ½, Semantics of Natural Language II pp. 18 – 67
- Lewis D. (1979a) "Score keeping in a Language Game" in *Journal of Philosophical Logic* Vol. 8 No. 1 pp. 339 – 359
- Lewis, D. (1979b) "Attitudes *De Dicto* and *De Se*" in *Philosophical Review* Vol. 88, pp. 513–543.
- Lycan (2005) "Names" in *The Blackwell Guide to the Philosophy of Language* (eds) Devitt and Hanley Malden: Blackwell Publishing
- Macfarlane (2007) "Semantic Minimalism and Nonindexical Contextualism" in (eds) Preyer and Peter *Context-Sensitivity and Semantic Minimalism: New Essays on Semantic and Pragmatics* Oxford: Oxford University Press

- MacFarlane (2009) "Nonindexical Contextualism" in *Synthese* Vol. 166, No. 2 pp. 231-250
- MacFarlane (2014) *Assessment Sensitivity: Relative Truth and its Applications* Oxford: Oxford University Press
- Magidor (2015) "The Myth of the De Se" in *Philosophical Perspectives* Vol. 29 pp. 249 – 283
- Margalit A. (1979) "Open Texture" in *Meaning and Use* (ed.) A. Margalit Dordrecht: D. Reidel Publishing company
- Marks (2017) *Sensory Blending: On Synaesthesia and Related Phenomena* Oxford: Oxford University Press
- Marr (1982/2010) *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information* London: The MIT Press
- McDowell (1984a) "Wittgenstein on Following a Rule" in *Synthese* Vol. 58, No. 3 Essays on Wittgenstein's Later Philosophy pp. 325-363
- McDowell (1984b) "De Re Senses" in *The Philosophical Quarterly* Vol. 34, No. 136 Special Issue: Frege, pp. 283 – 294
- McGinn (2015) "Evans on Understanding Demonstratives" in *The Philosophy of Language: The Classics Explained* accessed online: MIT Press
- McGurk and MacDonald (1976) "Hearing Lips and Seeing Voices" in *Nature* Vol. 264 No. 23 pp. 746 – 748
- Menno Lievers (2005) "The Structure of Thoughts" in (eds) Werning, Machery and Schurz *Foundational issues: The Compositionality of meaning and Content* Berlin: De Gruyter
- Millikan (1990) "The Myth of the Essential Indexical" in *Nous* 24 (5): 723-734
- Morris, M. (2007) *An Introduction to the Philosophy of Language* Cambridge, Cambridge University Press
- Müller (2008) "Representation in Digital Systems" in (eds) Briggie et al *Current Issues in Computing and Philosophy* IOS Press
- Neander (2017) "Jerry Fodor and Zenon Pylyshyn: *Minds Without Meanings: An Essay on the Content of Concepts*" in *Philosophical Review* Vol. 126 No. 3 pp. 410 – 417
- Newell A. (1980) "Physical Symbol Systems" in *Cognitive Science* Vol. 4 pp. 135 - 183
- Ninan, D. (2010) 'De Se Attitudes: Ascription and Communication'. *Philosophy Compass* 5: 551–567
- Ninan (2015) "On Recanati's Mental Files" in *Inquiry* Vol. 58, No. 4 pp. 368 - 377
- Ninan (2016) "About Oneself: De Se Thought and Communication" in (eds.) Garcia-Carpintero and Torre *About Oneself: De Se Attitudes and Communication* Oxford: Oxford University Press
- Ninan (2016b) "Assessment Sensitivity: Relative Truth and Its Applications" in *Philosophical Review* Vol. 125, No. 3 pp. 439 – 447
- Pagin and Westerståhl (2011) "Compositionality" in (eds) von Stechow, Maienborn, Portner *Semantics: An International Handbook of natural Handbook of Natural Language Meaning* Berlin:

Mouton de Gruyter Accessed from:

<https://stockholmuniversity.app.box.com/s/y0bptw0djyau0ijz3xtlwm6rh8vdsuq>

Papineau (2006) "Phenomenal and Perceptual Concepts" in (eds) Alter and Walter *Phenomenal Concepts and Phenomenal Knowledge: New Essays on Consciousness and Physicalism* pp. 111-44 New York: Oxford University Press

Partee (1973) "Some Structural Analogies Between Tenses and Pronouns in English" in *Journal of Philosophy* Vol. 70 No. 18 pp. 601 – 609

Partee (1979) "Semantics – Mathematics or Psychology" in *Semantics from Different Points of View* pp. 1 - 14

Partee et al. (1990) *Mathematical Methods in Linguistics* Dordrecht: Kluwer Academic Publishers

Peacocke (1981) "Demonstrative Thought and Psychological Explanation" in *Synthese* Vol. 49, No. 2 Demonstrative and Indexical Reference, Part II, pp. 187 – 217

Peacocke (1999) *Being Known* Oxford: Oxford University Press

Peacocke (2014) *The Mirror of the World: Subjects, Consciousness, and Self-Consciousness* Oxford: Oxford University Press

Perry (1977) "Frege on Demonstratives" in *The Philosophical Review* Vol. 86, No. 4 pp. 474-497

Perry (1986) "Thought without Representation" in *Proceedings of the Aristotelean Society, Supplementary Volumes* vol. 60 pp. 137-151

Perry (1979/2010) "The essential Indexical" in (ed.) A.P. Martinich *The Philosophy of Language* Oxford: Oxford University Press

Perry (2017) "Indexicals and Undindexicals" in (eds.) Ponte and Korta *Reference and Representation in Thought and Language* Oxford: Oxford University Press

Prinz (2006) "Is the Mind Really Modular?" in (ed) Stainton *Contemporary Debates in Cognitive Science* Oxford: Wiley-Blackwell

Jacob, Pierre, "Intentionality", *The Stanford Encyclopedia of Philosophy* (Winter 2014 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2014/entries/intentionality/>>.

Jary (2016) "Rethinking Explicit Utterance Content" in *Journal of Pragmatics* vol. 102 pp. 24-37.

Pietroski P. (2008) "Minimalist Meaning, Internalist Interpretation" in *Biolinguistics* vol. 2.4 pp. 317-341

Pietroski (2018) *Conjoining Meanings: Semantics Without Truth Values* Oxford: Oxford University Press

Pinker, S. (1994) *The Language Instinct: How the Mind Creates Language* New York, Harper Perennial

Pinker S. (2008) *The Stuff of Thought: Language as a Window into Human Nature* London: Penguin Books

Pollock (1990) "Understanding the Language of Thought" in *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* Vol.58, No. ½, Symposium on Thought and Reference pp. 95-120.

- Portner (2005) *What is Meaning? Fundamentals of Formal Semantics* Malden: Blackwell publishing
- Putnam (2002) "Travis on Meaning, Thought and the Ways the World is" in *The Philosophical Quarterly* Vol. 52, No. 206, pp. 96-106
- Raftopoulos (2001) "Is Perception Informationally Encapsulated? The Issue of the Theory-Ladenness of Perception" in *Cognitive Science* Vol. 25 pp. 423 – 451
- Recanati (1997) "Can We Believe What We Do Not Understand?" in *Mind and Language* Vol. 12 No. 1 pp. 84 -100
- Recanati (2004) *Literal Meaning* Cambridge: Cambridge University Press
- Recanati (2012) *Mental Files* Oxford: Oxford University Press
- Recanati (2015) "Replies" in *Inquiry* Vol. 58, No. 4 pp. 408 – 437
- Recanati (2016) *Mental Files in Flux* Oxford: Oxford University Press
- Recanati (2017) "Cognitive Dynamics: A New Look at an Old Problem" in (eds.) Ponte and Korta *Reference and Representation in Thought and Language* Oxford: Oxford University Press
- Robbins, Philip, "Modularity of Mind", *The Stanford Encyclopedia of Philosophy* (Winter 2017 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/win2017/entries/modularity-mind/>](https://plato.stanford.edu/archives/win2017/entries/modularity-mind/).
- Rovane (2013) "The Larger Philosophical Significance of Holism" in (eds) Lepore and Ludwig A *Companion to Donald Davidson* Chichester: John Wiley and Sons Inc.
- Rowlands (2010) *The New Science of the Mind: From extended Mind to Embodied Phenomenology* Cambridge: The MIT Press
- Rescorla (2017) "The Computational Theory of Mind" in *The Stanford Encyclopedia of Philosophy* (Spring 2017 Edition), (ed.) Zalta URL = [<https://plato.stanford.edu/archives/spr2017/entries/computational-mind/>](https://plato.stanford.edu/archives/spr2017/entries/computational-mind/).
- Rescorla, Michael, "Convention", *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/sum2017/entries/convention/>](https://plato.stanford.edu/archives/sum2017/entries/convention/).
- Salmon (1998) "Nonexistence" in *Noûs* Vol. 32, No. 3 pp. 277 - 319
- Samuels R. (2010) "Classical Computationalism and the Many Problems of Cognitive Relevance" in *Studies in History and Philosophy of Science* Vol. 41 pp. 280-293
- Sarnecki (2007) "Developmental Objections to Evolutionary Modularity" in *Biology and Philosophy* Vol. 22 pp. 529 – 546
- Scott-Phillips (2015) *Speaking Our Minds: Why Human Communication is Different, and How Language Evolved to Make it Special* Basingstoke: Palgrave MacMillan
- Schaffer and Szabo (2013) "Epistemic comparativism: a contextualist semantics for knowledge ascriptions" in [*Philosophical Studies* \(2\):1-53](#)
- Schmitt (2003) "The Production of Anaphoric Pronouns" in (eds.) Rickheit, Herrmann, Destsch *Psycholinguistics: An International Handbook* Berlin: De Gruyter Mouton
- Schneider (2009) "The Nature of Symbols in the Language of Thought" in *Mind and Language* Vol. 24, No. 5, pp. 523-553

- Searle (1975/2010) "Indirect Speech Acts" in (ed.) Matrini *The Philosophy of Language* 5th edition, Oxford: Oxford University Press
- Searle (1979) "Literal Meaning" in *Erkenntnis* Vol. 13 No. 1 pp. 207 – 224
- Searle (1980) "The Background of Meaning" in *Speech Act Theory and Pragmatics*
- Searle, J. (2002) "Can Computers Think?" in (ed.) D. Chalmers *Philosophy of Mind: Classical and Contemporary Readings*. New York, Oxford University Press
- Sennet, Adam, "Ambiguity", *The Stanford Encyclopedia of Philosophy* (Spring 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/spr2016/entries/ambiguity/>>.
- Siegel, Susanna, "The Contents of Perception", *The Stanford Encyclopedia of Philosophy* (Winter 2016 Edition), Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/win2016/entries/perception-contents/>>.
- Simms, K. (2009) "Language of Thought" in (ed.) C. R. S. Chapman *Key Ideas in Linguistics and the Philosophy of Language* Edinburgh, Edinburgh University Press
- Simon, H. (1996) "Machine as Mind" in (ed.) A. C. P. Millican *Machines and Thought: The Legacy of Alan Turing* Oxford, Clarendon Press. 1.
- Smith (2005) "Cognition" in (ed.) Honderich *The Oxford Companion to Philosophy* 2nd edition, Oxford: Oxford University Press
- Solt (2018) "Multidimensionality, Subjectivity and Scales: Experimental Evidence" in (eds) Castroviejo, McNally and Sassoon *The Semantics of Gradability, Vagueness, and Scale Structure: Experimental Perspectives* online copy: Springer
- Sperber D., Wilson D. (1995) *Relevance: Communication and Cognition*, Malden: Wiley-Blackwell
- Sprevak (2010) "Computation, Individuation, and the Received View on Representation" in *Studies in History and Philosophy of Science* Vol. 41, pp. 260-270. Downloaded from his website.
- Stalnaker (1976) "Propositions." Alfred MacKay and Daniel Merrill (eds.), *Issues in the Philosophy of Language*, New Haven and London: Yale U. Press, 1976, 79-91. With response by Laurence Powers
- Stalnaker (1981) "Indexical Belief" in *Synthese* Vol. 49, No. 1 pp. 129-151
- Stalnaker (2014) *Context* Oxford: Oxford University Press
- Stanley (2000) "Context and Logical Form" in *Linguistics and Philosophy* Vol. 23, No. 4 pp. 391 - 434
- Stern (2003) "Review: Unshadowed Thought: Representation in Thought and Language" in *Mind* 112 (448): 805-812
- Sternberg (2011) "Modular Processes in Mind and Brain" in *Cognitive Neuropsychology* Vol. 28: 3-4 pp. 156-208
- Stojanovic and Fernandez (2015) "Mental Files, Blown Up by Indexed Files" in *Inquiry* Vol. 58. No. 4 pp. 393 – 407
- Stokes (2013) "Cognitive Penetrability of Perception" in *Philosophy Compass* Vol. 8 No. 7 pp. 646 - 663

- Strawson (2003) "Mental Ballistics or the Involuntariness of Spontaneity" in *Proceedings of the Aristotelean Society*, New Series, Vol. 103, pp. 227 – 256
- Symons and Calvo (2014) *The Architecture of Cognition: Rethinking Fodor and Pylyshlyns Systematicity Challenge* Cambridge: The MIT Press
- Szabo, Z. G. (2013) "Compositionality" *The Stanford Encyclopaedia of Philosophy* (fall 2013 edition) (ed.) Edward N. Zalta
- Todd and Dieckmann (2004/2014) "Heuristics for Ordering Cue Search in Decision Making" originally in *Advances in Neural Information Processing* accessed from <https://www.researchgate.net/publication/221620559>
- Todd et al (2005) "Evolutionary Cognitive Psychology" in (ed.) Buss *The Handbook of Evolutionary Psychology* Hoboken: John Wiley and Sons, Incorporated
- Travis (1989) *The Uses of Sense: Wittgenstein's Philosophy of Language* Oxford: Oxford University Press
- Travis (1997) "Pragmatics" in (eds) Hale and Wright *A Companion to the Philosophy of Language* Oxford: Blackwell Publishers
- Travis (2000) *Unshadowed Thought: Representation in Thought and Language* Cambridge: Harvard University Press
- Travis (2006) *Thought's Footing: A Theme in Wittgenstein's Philosophical Investigations* Oxford: Clarendon Press
- Travis (2006a) "Insensitive Semantics" in *Mind and Language* Vol. 21, no1 pp. 39-49
- Travis (2017) "Views of my Fellows Thinking (Counting Thoughts)" in *Dialectica* Vol. 71, No. 3 pp. 337 – 378
- Travis (2018) "What Structure Lurks in the Minds of Men? Frege vs. Fodor" in *The Philosophy of Charles Travis: Language, Thought and Perception* Oxford: Oxford University Press
- Van Inwagen (1985) "Reviewed Work(s): The First Person: An Essay on Reference and Intentionality, by Roderick M. Chisholm" in *Nous* Vol. 19 no. 1 pp. 122-129
- Wheeler (2008) "Cognition in Context: Phenomenology, Situated Robotics and the Frame Problem" in *International Journal of Philosophical Studies* Vol 16. No. 3 pp. 323-349
- Wilson and Carston (2007) "A Unitary Approach to Lexical Pragmatics: Relevance, Inference and Ad Hoc Concepts" in Noel Burton-Roberts (ed.) *Pragmatics* Palgrave -Macmillan accessed from Open Access
- Wittgenstein, L. (1952) *Philosophical Investigations*, Wiley-Blackwell
- Wright (1986) "Rule Following, Meaning, and Constructivism" in (ed.) Travis *Meaning and Interpretation* Oxford: Basil Blackwell Ltd.
- Wright (1989) "Wittgenstein's Rule-Following considerations and the Central Project of Theoretical Linguistics" in (ed.) A. George *Reflections on Chomsky* Oxford: Basil Blackwell

Wright (1989/2002) "meaning and Intention as Judgement Dependent" in (eds.) Miller and Wright *Rule-Following and Meaning* Chesham: Acumen

Wright (2001) *Rails to Infinity: Essays on Themes from Wittgenstein's Philosophical Investigations* Cambridge: Harvard University Press

Wright (2007) "Rule-Following without Reasons: Wittgenstein's Quietism and the Constitutive Question" in *Ratio (New Series)* XX 4 Dec. 2007