

For and Against Objective Meaning

Philosophy Bachelor Thesis
15 ECTS

Jan Kanis

Saturday 8 October, 2011

supervisor:
Albert Visser

Contents

1	Introduction	2
2	Positions and Argumentation	4
2.1	Historical Background	4
2.2	Problem Description	5
2.3	Literalist Approaches	7
2.3.1	Traditional Literalism	7
2.3.2	Terminology: Pragmatic Processes	8
2.3.3	Minimalism	10
2.3.4	Indexicalism	14
2.3.5	Syncretism	19
2.4	Contextualism	23
2.4.1	Quasi-Contextualism	23
2.4.2	Modulation	23
2.4.3	Four Approaches to Contextualism	25
2.4.4	Meaning Eliminativism	26
2.5	Recanati's Position	30
2.5.1	Recanati's proposed model	32
3	Critical Evaluation	35
3.1	Other Approaches in Minimalism	36
3.2	More General Lines of Argumentation	37
3.3	Formal Rigour and Future Research	38

Chapter 1

Introduction

If I tell you that “Paul ran to the edge of the cliff and jumped”, your reaction might be “poor Paul!”. Not because he ran to the cliff, running is healthy and cliffs provide great views, but because he jumped. But if I were to speak just the sentence “Paul jumped”, where you would not have the former association of jumping *into an abyss*, you most likely would not at all react that way. But I did not actually tell you that Paul jumped *over* the cliff, I just said that he jumped. He might just be jumping up and down in the air on top of the cliff.

If this sentence were uttered in a normal conversation, the speaker would probably intend to convey the intuitive meaning, and not that Paul jumped up and down in the air. But does such a sentence only have this meaning because the speaker intends it, or can we say that objectively speaking, this sentence uttered in a normal conversation has the meaning we intuitively ascribe it.

According to the dominant position among linguists and philosophers of language, natural language sentences have a definite meaning (such as truth-conditions) independent of what the speaker means by the sentence. Many theories have been developed to try and capture this speaker-independent meaning. These theories have been able to describe a number of aspects of natural language, but there always remain exceptions which the theories do not handle correctly.

However, there is also another school. This school thinks that the goal of finding a theory of speaker-independent meaning is misguided because such a meaning does not exist. The reason, according to this school, is because the mainstream holds assumptions that are not warranted.

In order to get a more complete view of what these two schools of thought are, and what other positions exist regarding the question of speaker-independent meaning, I have studied “Literal Meaning”¹, a book by François Recanati. In this book, Recanati gives an overview of what positions there are regarding the above question. He also describes his own position in the debate, and de-

¹F. Recanati, ‘Literal Meaning’, Cambridge University Press, 2004. Further references to ‘Recanati’ will be to this book.

tails a lot of arguments that have been made for and against all these different positions.

In this thesis I will give an overview of the different positions that Recanati describes, and the arguments for and against them. I will also give the description and argumentation of Recanati's own position. This will constitute chapter 2. I only give a description of the positions regarding the question of speaker-independent meaning and the arguments that have been made. This is in fact what most of Recanati's book is about, but Recanati also makes a few detours regarding for example what we are to understand under "literal meaning". I did not include these in my summary. In chapter 3 I will give my own view on Recanati's argumentation, and I will give some future research suggestions that may help to come closer to a resolution of the debate.

While Recanati does a great job in providing the outlines of all the available positions and arguments, this only becomes apparent after reading the book for the second time. Unfortunately the book is not structured very well and it is very difficult to get an understanding of what Recanati describes and how all the parts fit together on the first reading. Recanati only gives the general overview of the available positions halfway through the book, and he has a strong tendency to mix descriptions of positions with the description and argumentation of his own position. Recanati describes most of his own position well before a general overview of all available positions is complete, meaning that at that point it is not yet possible to put his position into context.

In this thesis, I have tried to restructure the information Recanati gives in such a way that it is accessible on the first reading for someone who has the necessary background knowledge. I therefore start with a problem description and a broad view of the available positions. Terms and definitions are introduced as necessary, and I have made an effort to make sure all terms are defined before they are used.

Chapter 2

Positions and Argumentation

2.1 Historical Background

In the philosophy of language, there used to be two opposing camps: that of the Ideal Language Philosophers, and that of the Ordinary Language Philosophers. The Ideal Language camp was mainly concerned with formal semantics, analysing natural language as a system of formal logic. Their main focus was on studying what the truth conditions of expressions and sentences are. This camp included people like Frege, Carnap, Tarski and Russell². On the other hand, the ordinary language philosophers thought that analysing natural language as if it were a formal logic system was not the right way forward. Instead they advocated to study the use of language in everyday human interaction, as that would reveal what the meaning of an utterance of speech was. In this camp, people like Austin, Strawson, Grice³, and the late Wittgenstein were to be found.

Nowadays, this distinction is not as prevalent anymore, as both camps recognise each others achievements. The two views are now more commonly known as 'semantics' and 'pragmatics', and are seen as complementary fields, where semantics is concerned with meaning and truth conditions, and pragmatics with the use of language. But there is still a distinction remaining. Recanati discerns two positions, which he calls 'literalism' and 'contextualism'. According to literalism, the dominant view, we may legitimately ascribe truth-conditional content to natural language *sentences*, independent of what the speaker means. The other view, contextualism, holds that *speech acts* are the primary bearers of content and that one can not meaningfully talk about the meaning of a sentence independent of the speech act in which it was used. Since a specific speaker is always involved in a speech act, a speaker-independent objective meaning of a sentence does not exist.

²These pioneers were not originally concerned with natural language, but their methods were applied to natural language by their disciples such as Montague and Davidson.

³Although Grice is a special case, who has said he has one foot in each camp. (P. Grice, 'Retrospective Epilogue', in his *Studies in the Way of Words* (Harvard University Press, 1989), p. 372, cited in Recanati p. 1)

2.2 Problem Description

In the previous section, I have used the word 'meaning'. Since many of the different theories have different ideas of what a 'meaning' is, and many of them recognise several different kinds of meanings, let us start by defining the terms.⁴

There are lots of different ways to define what a 'meaning' of a word or sentence is, and to make matters worse, most theories on how language works use slightly or not-so-slightly different definitions, and often use several different types of 'meaning'.

Let us have a look at an example:

"I am French"

If I say this sentence, its meaning is that I am French. If you say it, it means you are French. Clearly, even though we utter the same words, the meaning is different. But there must also be some kind of meaning (where I use 'meaning' in the broadest sense) that can be ascribed to the sentence itself apart from who speaks it. If the sentence in isolation had no meaning whatsoever, then it wouldn't matter which words we actually use, and we wouldn't be able to predict accurately what it will mean when someone else utters these words. In fact, we can say that the sentence meaning of this sentence is that whoever utters it is French. Note that this meaning is not propositional: we cannot decide its truth value until who utters it is filled in.

This, then, brings us to the first distinction: that between the sentence- or linguistic meaning, and the speaker's meaning, or what is said by uttering it.

Linguistic- or Sentence Meaning:

The meaning a sentence has *qua* sentence-type. This kind of meaning is not in any way dependent on the context in which a sentence is used. It is a property of just the sentence on its own. The linguistic meaning is (according to most theories) not a propositional meaning, it is not something that can be true or false. According to some theories, the linguistic meaning can be represented as a formula of first-order logic with free variables. According to other theories, all that exists are the lexical meanings of words and the composition rules of the language, but the actual composition cannot be done without also including context.

What is Said:

The actual proposition (in case of assertive sentences) stated by whoever utters the sentence. This corresponds most closely to one's intuitive notion of the meaning of a sentence. All current positions on the literalism vs. contextualism debate agree that some amount of contextual influence is necessary, for example to resolve indexicals and demonstratives.

⁴I have also used the word 'content' without defining it. It is basically synonymous with 'meaning', and in the rest of this text I will use that.

This type of meaning is sometimes also known as ‘speaker’s meaning’, and in contextualist theories also as ‘contextualised meaning’ or ‘contextualised sense’.

We also need to make a second distinction. The sentence “*I am French*” can convey a lot more meaning than just that its utterer is French, if used in the right context. For example, if I ask someone if he can cook, and he answers by saying “*I am French*”, that clearly provides an affirmative answer even though that is not strictly what was said. In fact, if this same sentence got used in different contexts, it could convey lots of different meanings that may have nothing to do with being of French origin. These kinds of information conveyed by sentences are known as ‘implicatures’⁵. Unlike what is said, what is implied by a sentence is not bounded by the linguistic meaning of a sentence. If the context in which a sentence is uttered is changed enough, it can imply just about anything.

Implicature:

The implicature of an utterance is any kind of information that is conveyed beyond what is said. An utterance’s implicatures are not restricted by the words used in the utterance, unlike what is said by the utterance. If the context of an utterance changes enough, the utterance can imply just about anything.

Putting these three types of meaning together, we arrive at the following triad:

sentence meaning
vs
what is said
vs
what is implicated

These two distinctions, between linguistic meaning and what is said on the one hand, and what is said and what is implied on the other hand, have the advantage of being generally accepted as important to make by most linguists⁶. But that does not mean that they agree on what exactly each type of meaning encompasses.

The debate between literalism and contextualism is mainly concerned with the ‘what is said’ level. According to contextualism, ‘what is said’ must necessarily include a lot of contextual information, while according to literalists, ‘what is said’ includes a very limited amount of contextual information—or even none at all⁷—and that information is strictly controlled by the rules of the language.

⁵‘Implicate’ is a specifically linguistic term that encompasses the meaning of ‘imply’ in the specific sense of what a sentence implies. I will sometimes use ‘imply’ as a synonym for ‘implicate’.

⁶Or at least by those involved in the current debate.

⁷At that point the distinction between linguistic meaning of a sentence and what is said by

2.3 Literalist Approaches

In the following sections, I will introduce the positions in the debate, starting with those on the literalist side of the spectrum. Additional terminology will be introduced as necessary. For the benefit of the reader, here are the positions I will discuss:

- Traditional Literalism
- Minimalism
- Indexicalism
- Syncretism
- Quasi-Contextualism
- Contextualism

2.3.1 Traditional Literalism

The first position is traditional literalism. It is the oldest position, and actually is not so much a position in the literalism vs. contextualism debate, but the position of the ideal language philosophers in their debate with ordinary language philosophy. As such, this position has largely been abandoned by current philosophers as they began to integrate the ideas of ordinary language philosophy, but since this position is in many ways the ancestor of current literalist positions I will still give a very short description of what it entails.

The ideas resulted from the development of formal logical languages, in the tradition of Russell.⁸ Traditional literalists thought that the meaning of any natural language sentence ought to be representable as a formula of first order logic. So, for example, the sentence “*The president of America is bald*” would become

$$\exists x.(\text{PresOf}(\text{America}, x) \wedge \forall y.(\text{PresOf}(\text{America}, y) \rightarrow (x = y)) \wedge \text{Bald}(x))$$
⁹

(Where PresOf is an abbreviation of President Of.) With this tradition the idea was born that the meaning of a sentence equates with its truth conditions.

uttering a sentence disappears, so—as an exception to what I said before—this most extreme form of literalism does not recognise that distinction at all. But see the upcoming section 2.3.1.

⁸Russell himself was a logician, and his aim was not in fact to create an accurate description of how language works. However, other philosophers in this tradition built upon his work and applied it to natural language.

⁹The “ $\exists x.(\text{PresOf}(\text{America}, x) \wedge \forall y.(\text{PresOf}(\text{America}, y) \rightarrow (x = y)) \wedge \dots)$ ” part basically says that there is only one president of America. It represents the translation of the definite article “*the*”.

This approach has some difficulties when translating indexicals and demonstratives, because first order logic does not have a construct for them.¹⁰ The traditional literalists' response was that while some sentences might contain such 'difficult' words that referenced the environment, one could always rewrite a sentence containing such words into a sentence that did not, but instead used definite descriptions. So, for example, a sentence such as "*He is bald!*" (uttered while pointing to the president of America appearing on TV) would translate into something like "*The person appearing on the television on BBC1 on Saturday 8 October, 2011, 1:20pm is bald!*". Such a sentence that does not contain references to its environment is called an eternal sentence.

In the debate between the ideal language and ordinary language philosophies it has been argued that many mechanisms, such as reference and quantification, are fundamentally context-sensitive, and that it is therefore impossible in general to transform every possible context dependent sentence into an eternal one. The only real rule for demonstratives and similar constructs is that they refer to *what the speaker intends them to refer to*, and that is not something that can be resolved without appeal to the wider context.

These arguments have generally been accepted and traditional literalism is now by and large abandoned as a position. Therefore I will not go deeper into the specific arguments.

2.3.2 Terminology: Pragmatic Processes

The next position we are going to discuss is minimalism. Minimalism is a direct descendant of traditional literalism, staying close to its roots but incorporating more recent findings. Minimalism accepts that there are context dependencies in language, so it no longer insists that every context dependent sentence can be reformulated into an eternal sentence, however it tries to keep the amount of contextual information that is used to formulate what is said to a minimum.

To see how it does that, we first have to take a closer look at the processes that incorporate contextual information into the sentence's meaning.

In most theories that are relevant to us, a pragmatic process can be seen as a process that takes a meaning of some kind as input, together with a context or some specific pieces of information from the context, combines them, and

¹⁰There are other difficulties as well, such as how to translate "*The king of France is bald!*". The natural translation would be

$$\exists x.(\text{KingOf}(\text{France}, x) \wedge \forall y.(\text{KingOf}(\text{France}, y) \rightarrow (x = y)) \wedge \text{Bald}(x))$$

however that formula is false not because the king of France fails to be bald, but because the king of France does not exist. The Russellian approach leaves no room for things like presuppositions and other properties that do not map directly to truth or falsehood. Other approaches do try to solve these problems by for instance changing the formal language to include presuppositions. However, those approaches are not immediately of interest to the current discussion.

produces another meaning, usually of a different kind. So, one could say that there are pragmatic processes that take the linguistic meaning of an uttered sentence as input, and produce what is said, and other pragmatic processes that take what is said as input and produce what is implied. However, as we shall see, this very much oversimplifies things, and it does not take into account that the different theories sometimes have very different views on how exactly the different kinds of meaning are constructed. It is, however a useful picture to keep in the back of our minds for now.

When we look back at our example of *"I am French"*, the first and most obvious instance of context dependence is the indexical *"I"*. Clearly, indexicals and similar constructs need to be resolved if we want to be able to determine the truth conditions of a sentence. On the other hand, there is the contextual enrichment that can (in certain contexts) make this sentence imply *"I am a good cook"*. This additional information, while an important part of the conversation, is not strictly necessary in order to determine the truth or falsehood of the sentence spoken. In addition to one of these instances of pragmatic processes being mandatory while the other is not, we can also observe that the first process is driven by the presence of certain words in the sentence, i.e. the indexical, while the second is driven by features of the context, in this case that there is a question to be answered.

This, then, brings us to the first distinction to be made concerning pragmatic processes:

Mandatory Pragmatic Processes a.k.a. Saturation:

Processes that happen bottom-up, controlled by the structure of the sentence. These types of processes are mandatory because without them the meaning being built cannot constitute a full proposition. In minimalism, this type of processes is called 'saturation'.

Optional Pragmatic Processes:

Processes that are driven by the context, i.e. top-down. These processes are not mandatory, the sentence also has a propositional meaning without applying these processes.

There is also a second property of pragmatic processes that should be distinguished. Implicatures often appear to be derived from a previous proposition, typically from what is said. Other contextual influences, such as the resolution of indexicals, do not require a propositional meaning as input. Recanati distinguishes this aspect of pragmatic processes as being primary or secondary.

Primary Pragmatic Processes:

A pragmatic process that does not require propositional input. The input can be what is called pre-propositional.

Secondary Pragmatic Processes:

Pragmatic processes that require a fully formed proposition as input, and give an output that is based on it.

According to Recanati, normal interpreters without linguistic training usually are not aware of the pre-propositional input into primary pragmatic processes, only in their output.

While the properties of mandatory and primary on the one hand, and optional/secondary on the other hand often coincide (such as with the “*I am French*” example), according to Recanati they should be distinguished as independent and orthogonal. Not all theories make this distinction, however. For many theories only the mandatory/optional aspect is important, and they do not deal with the primary/secondary dimension at all. Minimalism is, in fact, a theory that does not care about this latter dimension. Recanati’s own model, as we shall see, does.

2.3.3 Minimalism

The basic premise of minimalism is that the linguistic meaning of a sentence can be represented as a formula of some first order logic¹¹, but, unlike traditional literalism, minimalism does not require that a context dependent sentence can be reformulated as an eternal sentence. For the logical form, this means that the linguistic meaning of a sentence is allowed to be an open formula with free variables.

Minimalism then asserts that what is said by a sentence is derived from its linguistic meaning by *saturation*. Saturation, as we saw above, is another term for what Recanati calls ‘mandatory pragmatic processes’, processes which must take place in order to create a truth-evaluable meaning from the pre-propositional linguistic meaning.

The main task of saturation consists of filling in open variables with items from the context. (But it may also include adding ‘unarticulated’ constituents to the sentence, if one assumes, as some philosophers do, that that is required in order to make the sentence truth-evaluable.) The result of saturation is the so-called *minimal proposition*:

minimal proposition:

The (according to minimalism) first truth-evaluable form of meaning that is derivable from a sentence. The minimal proposition is the result of applying saturation to the sentence’s linguistic meaning.

Let us take the sentence

¹¹First order predicate logic is one option, but approaches that use for example lambda calculus or Montague grammar can also be included under either minimalism or indexicalism (see section 2.3.4), depending on the details of what they do.

“He is tall”

as an example. This sentence could have a linguistic meaning of the form

$\text{Tall}(x)$ (with an additional constraint that x is male)

x is a free variable, so this formula does not have a truth value. If this sentence were used in a specific situation, the process of saturation would fill in any free variables, resulting in for example

$\text{Tall}(x)[\langle \text{John} \rangle / x] = \text{Tall}(\langle \text{John} \rangle)$ ¹²

as the minimal proposition (and thus also as what is said).

Saturation also happens for other constructs, such as genitives. In the phrase *“John’s painting”*, it needs to be determined if the relation is that of being the painter of the painting, or being the owner. The linguistic meaning of such a phrase would be something like *“the painting that bears a certain relation R to John”*. Other examples are parametric predicates such as *“small”* (something is always small compared to something else) or *“on the left of”*, and definite null instantiation such as in *“I noticed”*.¹³

Criticism

It is important to note that, while saturation is a pragmatic process, it is a process that is driven by the linguistic meaning, not one driven by the context. The linguistic meaning is generated without any kind of context being involved. Certain lexical and syntactic constructs lead to free variables becoming part of the linguistic meaning, and the context then only gets used to fill in those specific free variables. There is, in the minimalist framework, no way in which the context could influence other aspects of what is said by a sentence other than those for which a free variable exists.

This is a problem, because according to Recanati there *are* in fact other pragmatic processes which influence a sentence’s truth conditions but which are not mandatory, and therefore are not instances of saturation. He gives a number of categories of such processes. They can be divided into forms of free enrichment, loosening, and ‘semantic transfer’. This list is not meant to be exhaustive, but to show that saturation alone is not enough.

¹² At this point there is a slight complication concerning what ‘ $\langle \text{John} \rangle$ ’ is exactly. I will not go into that as that would open up a whole different can of worms. Allowing a logical constant to denote John entails that the result of saturation would be an eternal sentence, but minimalism does not use those. The easiest way out for now is to let ‘ $\langle \text{John} \rangle$ ’ be John himself, the actual person that is referred to.

¹³ *“I noticed”* contrasts with a sentence such as *“I ate”*. In the former case, it must be contextually provided *what* got noticed, while the latter case is equivalent to *“I ate something”*, where what is eaten is left unspecified.

Free Enrichment:

A sentence constituent is 'enriched', it is interpreted in a contextually more specific way, as if an additional unarticulated constituent was provided. For instance, the word "*rabbit*" can be interpreted differently in different contexts. In the sentence "*He wears rabbit*", "*rabbit*" will be interpreted as RABBIT FUR, while in "*He eats rabbit*", it will be RABBIT MEAT, and in the sentence "*After the accident, there was rabbit all over the highway*" it is RABBIT STUFF.

Free enrichment can also add additional restrictions to the application of a predicate. For example, there are certain conditions of application that form the definition of Table, and enrichment can add a further condition `In_the_living_room` if we are only concerned with those tables. This mechanism can also account for the contextual restriction of quantifier domains, for example if in "*All the books are on the table*" we are not talking about all the books in the world but about a particular set of books and a particular table.¹⁴

Loosening:

Loosening is the opposite of enrichment. Rather than interpreting a constituent more specifically than its regular context independent meaning, it is interpreted as if something more general was said. Take for example "*The ATM swallowed my credit card*". There is no actual swallowing going on since ATMs are not living organisms that can swallow, but the verb "*swallow*" is applied to what the ATM does anyway.

Semantic Transfer:

This is mostly a grab bag category of optional truth-conditionally relevant pragmatic processes that do not fall under one of the other categories. A concept is not enriched nor impoverished relative to its out of context interpretation, but it is changed in a way that is still related. An example of this is "*I am parked out back*", where by "*I*", I mean "*my car*".

According to minimalism, the results of these processes are not incorporated into the minimal proposition. Rather, they take part in constructing a sentence's implicatures. However, this can in many cases lead to minimal propositions that

¹⁴Recanati makes a further subdivision between two forms of free enrichment: specification and strengthening. I will not go into this further distinction, except for saying that the rabbit example is a case of the former while the latter is concerned with restricting applications of predicates.

The effects of free enrichment can often also be analysed as providing additional unarticulated constituents, instead of changes of the logical form or the creation of ad hoc concepts as happens in this description. Recanati handles this as if there is a single form of free enrichment which can be handled in different frameworks, either in terms of unarticulated constituents or in the way Recanati describes. The choice of framework is not currently relevant.

are very different from what is intuitively said by a sentence. Let us consider some more examples:¹⁵

1. *I've had breakfast.*
2. *You are not going to die.*
3. *The table is covered with books.*
4. *Everybody went to Paris.*
5. *John has three children.*

In all these cases, the process of saturation, without applying any optional pragmatic processes, implies that what is said by these sentences according to minimalism is not what one would intuitively ascribe to them. "*I've had breakfast*" would according to the minimalist position be true if the speaker *S* has had breakfast before the time of utterance of this sentence *t*. So strictly speaking, it would be true if the speaker has had breakfast twenty years ago and never since. That is clearly not what a speaker usually means with this sentence, for example if used as a reply to the question "*Would you like something to eat?*" In such a case the speaker would mean something along the lines of "*S has had breakfast on day d*", where *d* is today. This aspect of the speaker's meaning has to be construed as something external to what is said, i.e. as an implicature, in the same way that "*I am French*" can imply that the utterer of that sentence is a good cook. This is so because the minimal interpretation, to the effect that *S* has had breakfast some time in her life, is sufficient to make the utterance into a well formed proposition. There are no free variables that require to be filled in. In fact, in a different context the minimal proposition might actually be exactly what is meant by the speaker.

Something similar can happen in the second example. Imagine a child crying because of a minor cut and her mother uttering (2) in response. What she means is "*You are not going to die from that cut*", but according to minimalism, what is said by this sentence is that the child is not going to die *ever*, as if he or she were immortal. The extra element provided—the implicit reference to the cut—does not correspond to anything in the sentence, nor is it an unarticulated constituent whose provision is necessary to make the utterance into a full proposition. Like the previous example, we can easily imagine contexts in which what is meant by this sentence is only the minimal proposition.

Examples (3) and (4) can be analysed in the same way. In a standard Russellian analysis a definite description conveys uniqueness, so "*The table is covered with books*" is true iff there is one and only one table, and it is covered

¹⁵These example sentences are taken from Recanati. He has one more sentence, "*It's raining*", but I excluded that because I do not agree with Recanati's analysis and the point he tries to make is clear enough without it.

with books. If this is to make sense, we either need to focus on a restricted situation in which there actually is only one table, or we need to enrich the predicate *table* into *table_of_the_living_room* or something along those lines. Either way, it is arguable that the form of enrichment we use is not linguistically mandated, but only pragmatically required to make sense of the statement. So according to minimalism it cannot be part of what is said.

Example (4) has the same problem. Without restricting the meaning of “*everybody*”, that expression would refer to everyone in the world. While it would be possible that the utterer of (4) means that everybody in the world went to Paris, in most situations he or she would mean ‘everybody from a restricted group of people’, which again is a restriction that is not linguistically required and hence not part of the minimalist proposition.¹⁶

For (5), we need to think about what “*three*” means. A fairly standard view¹⁷ is that numerals like “*three*” mean “*at least three*”. This is true in sentences like “*If John has three children he can benefit from lower rates on public transport*”. In (5) however, the intuitive reading is “*John has exactly three children*”. According to minimalism, what the speaker says by uttering this sentence is that John has at least three children, and by invoking the maxim of quantity he has implied that John has no more than three children. However, anyone uttering such a sentence usually is not at all aware of having said something resembling “*at least three children*”, because the words “*at least*” were not ever uttered.

These sentences show that the minimal proposition as construed by minimalism is often not what we intuitively think we said, and neither is it something that is necessarily part of what is said. As it is, the minimal proposition starts to look very much like a purely academic exercise without any grounding in reality.

2.3.4 Indexicalism

Indexicalism is a more extreme form of minimalism, that has come up as a response to the contextualist criticisms against minimalism. Indexicalism follows the same approach in saying that a sentence gets converted into a logical formula with open variables, which then get filled in by saturation, leading to the minimal proposition (which is also what a sentence says according to indexicalism). The difference is, that in indexicalism there are much more free variables present.

¹⁶Recanati’s analysis of this example depends on “*everybody*” quantifying over all people in the world. In predicate logic, this is the normal way to interpret quantification, but I doubt if that should also be assumed for natural language. I think, given that the domain of quantification is one of those things that is highly context dependent, that it is more natural to model quantifiers as quantifying over a free variable, which is to be filled in from the context. If quantifiers are modelled in that way, this example is no longer a problem for minimalism because now the domain of quantification can be filled in by saturation.

¹⁷Larry Horn, “The Natural History of Negotiation”, Chicago University Press, 1989, pp. 205–16, referenced in Recanati p. 11 footnote 9

While in minimalism there's a rough correspondence between the words used in the sentence and the free variables in the resulting formula, in indexicalism this correspondence is lost. The criticisms of contextualists are countered by assuming free variables to account for the contextual dependencies that are missing in minimalism. In that way more contextual dependencies can be pulled into the minimal proposition by the process of saturation.

Indexicalism still holds on to the possibility of eternal sentences in principle, but believes that there are so many words or constructions that introduce free variables that in practise one rarely encounters them.

Like minimalism, indexicalism says that the minimal proposition states the truth conditions of a sentence, so contextual elements can only influence the truth condition if they are 'pulled in' to the sentence by saturation. Optional pragmatic processes still can not influence the sentence's truth condition.

The difference between indexicalism and minimalism is that in minimalism free variables only exist if they are provided by specific lexical elements or grammatical constructs, while in indexicalism this is not a necessity. According to indexicalism a free variable also exists for a specific contextual element if in every context where a sentence can be used, such an element has to be provided.¹⁸

This is summarised in the optionality criterion, which determines if a contextual element gets handled by saturation or by optional pragmatic processes:

Optionality Criterion:

A dependency on an element from the context is optional if the sentence can be used in a different context in which no such element is provided, yet the utterance still forms a complete proposition.

If an element fulfils this criterion, it is an optional contextual ingredient. The element is handled by optional pragmatic processes and can therefore not influence the sentence's truth conditions. But if an element does not fulfil this criterion, it is a mandatory contextual ingredient, and a free variable must be assumed to exist for it even if the words and the grammatical structure of the sentence do not indicate so. In that case the element is processed by saturation and is able to influence the sentence's truth conditions.

At this point we should look at another example as an illustration. Suppose we ask someone to open the door. The addressee must first of all identify the relevant door, but she must also determine the sense in which it is to be opened.

¹⁸This distinction is quite subtle, because one wonders what happens if a sentence is used in a context where not all mandatory contextual elements are provided: in that case, the sentence probably fails to be truth-evaluable. But this can also be explained in a minimalist framework by assuming that the lexical elements of the sentence in question have more variables that do not get filled in. In that case these unfilled variables are also present in the linguistic meaning, but now they are provided by lexical elements.

This lack of a clear boundary is, however, not a real problem. Minimalism and indexicalism are usually not in conflict with each other, rather one is an extension of the other. The positions should be seen as two points on a continuum with a smooth transition between them.

Besides doors and windows, eyes and wounds can be opened. If the addressee were to open the door by making an incision in it with a scalpel, the way one opens a wound, she would not have satisfied our request. But still, in a special context, we could imagine that this precisely is the way we expect the request to be fulfilled. Even if we were to provide the mode of opening explicitly, there will always be some other aspect of the door-opening that will be underspecified.

For minimalism, this is a problem because the verb 'to open' is usually not considered to be indexical, and neither is there a special grammatical construct present to warrant indexicality. The indexicalist would answer as follows: Like all verbs, the verb 'to open' is associated with a complex frame¹⁹ involving a number of argument roles: a location playing the role of inside, another location operating as outside, a boundary separating the two, a moving_object which is to pass from inside to outside (or the other way around), an obstacle preventing the passage of the moving object, an agent liable to free the passage by means of action on the obstacle, an instrument serving to accomplish the action, and so on and so forth. In the context, each of these variables²⁰ must be given a value. In the case of opening a door, the inside would be one room, outside the next, moving_object the person who wants to go to the other room, the boundary the door, and so on.

Taking this approach, the indexicalist appears to be able to answer any example that contextualists have identified as problematic for minimalism. Any expression, such as most verbs, can be associated with a frame with argument roles that provides a lot more free variables to the sentence's linguistic meaning than would be available in a minimalistic analysis.

But contextualists like Recanati are not fully satisfied with this answer to their criticisms. A problem remains in deciding whether a contextual dependency is optional or not. The optionality criterion pretends to give a clear and unambiguous answer to this question, but in fact there are cases where it is very hard to clearly decide whether the criterion applies or not.

Let us take as an example the sentence "*Most students are male*". This sentence can be uttered by for example a teacher in a context where "*most students*" refers to most students in his class. The sentence can also be uttered in a different context, without such a restriction, where it refers to all students, making a statement about students in general.

The sentence could be analysed so that the domain restriction is a pragmatic influence, provided by optional pragmatic processes, but it could also be analysed by saying that a free variable does exist for the domain, where for the second, general utterance we assign the domain of all students that exist to that variable. In order to decide what is the case, we will need to apply the optionality criterion.

¹⁹See Charles Fillmore, "Frame Semantics and the nature of Language" in *Annals of the New York Academy of Sciences* 280 (1976), pp. 20-32, and other papers by Fillmore

²⁰Each variable that is not optional according to the optionality criterion. Recanati does not explicitly apply the criterion to this example, but it probably should be.

If the domain restriction is optional, indexicalism prescribes us to use the former analysis, if it is not optional, we will use the latter.

But how are we to decide whether the domain restriction is optional or not? The first context clearly provides a contextual ingredient restricting the quantifier domain, but in the second context things are unclear. We are not aware of any assignment of free variables going on, but that does not mean there is not. Our brains could be assigning the set of all students to the quantified-over variable, without us being consciously aware of it. Since nothing is actually being restricted there is no reason why we would need to become consciously aware of such a process.

This leaves us with a problem: we are unable to decide if the optionality criterion applies or not. Our intuition does not give us an answer, and neither do the principles of indexicalism.²¹

Stanley's Binding Criterion

As an alternative to the optionality criterion, Jason Stanley introduced the Binding Criterion:

Binding Criterion:

A contextual ingredient in the interpretation of a sentence results from saturation if it can be 'bound', that is, if it can be made to vary with an operator in whose scope it lies.

Operators, specifically quantifiers, are part of the linguistic meaning of a sentence and can influence the interpretation of variables which lie in their scope. Therefore, if changing an operator also modifies the interpretation of a contextual ingredient, that is proof that the contextual ingredient was introduced into the propositional meaning to fill a free variable.

Some examples:

- *"In every room in John's house, every bottle is in the corner."*
- *"Whatever John does, most of the class falls asleep."*

In both of these examples the first quantifier ("*every*", "*whatever*") takes wide scope over the second quantifier ("*every*", "*most*"). In the first sentence, the scope of "*every bottle*" is 'every bottle within each room', in the second, in

²¹ We could, in order to disambiguate such cases, say that assigning the entire domain counts as not making an assignment, rendering the variable optional. This would not help the indexicalist's position since that would give us the wrong analysis for the example of "*Most students are male*" with "*most*" referring to the current class. Saying that in such cases the variable is mandatory does not help either, because then any quantified-over variable would be mandatory. Since there is no bound to the number of variables that could implicitly be quantified over, as with the 'open' example, this would lead to a possibly unbound number of variables in any sentence's linguistic meaning.

each case John does something, a different part of the class may fall asleep. If one were to change the first quantifier, for example to “*In some room. . .*” and “*For one thing that John does, . . .*”, the interpretation of the second quantifier would also change. In the first example we would only be talking about one single corner and every bottle in the house instead of every bottle in each room. In the second, we are not talking anymore about a different part of the class on each occasion. If the second halves of these sentences did not contain free variables, these differences could not occur, so in both of these cases we are dealing with saturation.

The binding criterion, according to Recanati, also fails to provide a firm foundation for indexicalism. In some cases it is too strong, also applying to cases that even Stanley agrees are pragmatic. For example in “*The policeman stopped the car*”, this can be by signalling the driver, or by pressing the brakes if he is driving it himself. We can construct the sentence “*However he did it, the policeman stopped the car*” meaning “*In whatever manner of stopping m he did it, the policeman stopped the car in manner m .*” The manner of stopping is bound by a prefixed operator, and therefore this would be saturation.

The same happens in “*Whenever his father cooks, John eats*”. Our intuitive meaning is that this means that John eats the food his father has cooked, which implies that there must exist a variable representing the food, that gets bound.

Let us dive a bit deeper into what is wrong with the binding criterion. The phrase “*whenever his father cooks*” operates on “*John eats*”. The latter is in itself a complete sentence which can also be uttered in isolation. Stanley’s argument relies on the presupposition that any variables that are present when “*John eats*” is used as an operand are therefore also present when this sentence is uttered in isolation. According to Recanati this need not be true. The variable that gets modified need not be part of the original predicate, but can be contributed by the modifying expression.

Recanati comes up with an alternative analysis, that interprets the modifying expression as a second order function that changes the arity of the original predicate. To do this he introduces the *Circ* operator. Applying *Circ* to *Eats*, we get the following:

$$\text{Circ}_{\text{location}}(\text{Eats}(x)) = \text{Eats_in}(x, l)$$

The whole sentence “*John eats in Paris*” gets interpreted as

$$\text{Circ}_{\text{location:Paris}}(\text{Eats}(\text{John})) = \text{Eats_in}(\text{John}, \text{Paris})^{22}$$

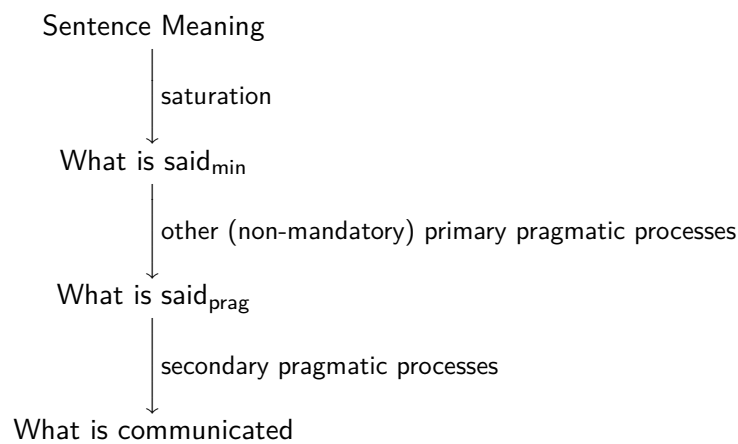
Recanati names Stanley’s assumption that any variable that is to be modified must be part of the original predicate, the *Binding Fallacy*. Stanley’s argument

²²This formal notation is a bit sketchy, but this is what Recanati uses. A more rigorous notation would be: $\text{Circ}_{\text{location}}(\text{Eats}) = \text{Eats_in}$ where *Eats* is of type $e \rightarrow t$ and *Eats_in* of type $e \rightarrow (e \rightarrow t)$. The interpretation of the full sentence then becomes $\text{Circ}_{\text{location:Paris}}(\text{Eats})(\text{John}) = (\lambda x. \text{Eats_in}(x, \text{Paris}))(\text{John}) = \text{Eats_in}(\text{John}, \text{Paris})$

is based on the binding fallacy and is therefore not convincing.

2.3.5 Syncretism

The next position on the spectrum is actually an in-between position between the literalist varieties and contextualism. Syncretism tries to unite the two camps, by proposing that there are two notions of what is said: one based on the ideas of minimalism, and one based on pragmatic notions. The result is a four layer model:



The ‘what is said_{min}’ level is the result of saturation, so in that sense it is the equivalent of the minimal proposition in minimalism and indexicalism. The ‘what is said_{prag}’ level is a form of meaning that includes other optional primary pragmatic processes, and so is able to include a lot of context-driven contextual influences. This is, as we shall see, very similar to the contextualist notion of what is said.

There are several authors who have proposed models like the one above²³, and there are also different ways in which the two levels can be construed. In this section we will have a look at several forms of syncretism, and see if they make true on their promise of unifying literalism and contextualism.

Literalist Interpretation

One popular way to define the two layers of what is said is to say that what is said_{min} is what *the sentence* says (with respect to the current context), while what is said_{prag} is what *the speaker* says by uttering that sentence. In this interpretation the what is said_{min}/what is said_{prag} distinction is closely related to the literalist construal of the semantics/pragmatics distinction. However, that distinction is in itself very questionable for a contextualist like Recanati.

²³ See for example N. Salmon, “The Pragmatic Fallacy”, in *Philosophical Studies* 63, 1991, pp. 83–97; S. Soames, *Beyond Rigidity*, p. 86; J. Berg, *In Defense of Direct Belief*, pp. 466–7.

In the semantic level of meaning, according to this literalist distinction, knowing a language is like knowing a theory by which one can deductively establish the set of truth conditions under which a sentence in that language holds. This interpretation of a sentence can be constructed mechanically by following the rules of the language. This therefore also holds for what is said_{min} under the literalist interpretation.

The pragmatic form of meaning, which corresponds to what is said_{prag}, is something totally different. It is not concerned primarily with language, but with human action. The pragmatic interpretation is defeasible: you can never be totally sure you have the right interpretation. Even if you have a lot of evidence for one, it is always possible that stronger counter-evidence is discovered later, which proves your original interpretation was wrong after all.

The problem with this model, according to Recanati, is that it is not always possible to construct a semantic interpretation independent of the pragmatic interpretation. A phrase like “*John’s car*” means something like “*the car that bears relation R to John*”, where R is a free variable. There are no rules in the language that determine how exactly R should be filled in, that is something that can only be done based on the wider context. In the same way, a word like “*here*” can refer to this room, this building, this country, etc. Although such an indexical provides a free variable in the linguistic meaning (according to the minimalist model), there are no language rules to determine what the correct value of that variable should be. Choosing the correct interpretation depends on the wider context, including the speaker’s intentions. Although minimalism greatly limits the influence context can have on what is said, saturation is not a mathematical process that objectively gives one unambiguous value to each free variable.

Recanati claims that such ambiguities can occur with many words and constructs, so that in general there is no such thing as what *the sentence* says, there is just a single notion of what is said, and that notion is pragmatic.

The Minimal Proposition as ‘Common Denominator’

A different interpretation of what is said_{min} put forward by Soames, is to say that what is said_{min} is abstracted from what is said_{prag}. It is the shared common denominator of all possibilities for what is said_{prag}. According to Recanati this could hold in the case of enrichment, but not in the case of semantic transfer or loosening. If in a bar, a waiter says: “*The ham sandwich left without paying*”, thereby referring to the ham sandwich orderer, he does not assert that the sandwich itself left without paying. If we agree that it must be part of what is said_{min} that the ham sandwich itself left, that cannot be abstracted from what is said_{prag}, because it is not part of it.

Reflexive Interpretation

A third interpretation of what is said_{min} is to interpret it as reflexive meaning. If someone utters the sentence *u*: “*I am French*”, the interpretation of that sentence is taken to be “*The utterer of *u* is French*” (where *u* stands for the specific utterance). Such an interpretation allows to avoid any direct references to contextual elements by always indirecting through what the ‘utterer of *u*’ meant or intended.

However, this interpretation is problematic when we want to derive what is said_{prag} from what is said_{min}. What is said_{prag} should still be in accordance with what we intuitively think we say. But if for example François utters the above sentence, what he intuitively thinks he has said is something like “*{François} is French*”, and not “*The utterer of this utterance is French*”. So the indirection through ‘the utterer of *u*’ that was introduced in what is said_{min} needs to be removed again in what is said_{prag}. But there are no recognised optional secondary pragmatic processes that would do such a job. This interpretation would therefore require a new kind of pragmatic process to be introduced, which would probably be quite complex in order to correctly do its job, and which would therefore be hard to justify. As far as I know no one has tried to propose such a thing.

A second problem that Recanati identifies is that this interpretation of what is said_{min} does not actually depend on saturation. All free variables that exist in the minimalist or indexicalist sentence meaning are replaced by expressions containing ‘the utterer of *u*’. With saturation gone, the reflexive interpretation is entirely determined by the language rules, and therefore collapses into the sentence meaning. Recanati concludes that while the reflexive proposition can be legitimately formed, it leads to a collapse of the four layer model. The result is a three layer model with the ‘what is said’ level being a pragmatic one.

I do not agree with Recanati on his second argument. While the reflexive interpretation replaces most free variables, it still has one left: the reference to the utterance, which obviously is a contextual element. This avoids a lot of the arguments against minimalism because a lot of cases that are problematic for saturation are removed. However it still leaves a what is said_{min} that is not a part of nor implied by what is said_{prag}, and it pushes all the real problems into the unspecified pragmatic process that translates one into the other. The reflexive proposition does not help us in understanding the utterance because for any mildly context dependent sentence the reflexive proposition would include so many ‘what the utterer of *u* intended’, ‘what the utterer of *u* pointed to’ etc. that we might as well say that the reflexive proposition is “*What the utterer of *u* meant is true*”. While I think that it is possible to construct the reflexive proposition, I do not think it is useful.

Other Problems

How we interpret one constituent of a sentence can be dependent on the interpretation of other constituents. And this dependency is not limited to one level, it can be recursive. In syncretism, saturation leads to what is said_{min}, and other optional pragmatic processes lead from what is said_{min} to what is said_{prag}. However, it is possible to create situations in which there is a circular dependency where mandatory pragmatic processes depend on the result of optional pragmatic processes.

Recanati shows this with an example. Suppose there is a fight between five warriors and five beasts. The beasts are a wolf, a lion, a hyena, a bear and an alligator, and each warrior has to battle one of them. Each warrior carries a shield with his coat of arms on it, one of them having a lion as its coat of arms. After the fight is over, the five beasts lie on the ground with swords through their bodies. Now assume that in this context the sentence *“Give me the lion’s sword—I want to have a look at it”* is uttered.

In this context, *“the lion”* can be interpreted as literally referring to the lion, or non-literally referring to the warrior who has a picture of a lion on his shield. The possessive relation can also be interpreted in different ways: *“X’s sword”* can mean *“the sword that X used in the fight”* or *“the sword now sticking through X’s body”*. In this context, both are a possibility depending on what the interpretation of *“the lion”* is. Now suppose that what the speaker actually means is to be given the sword that the warrior with the picture of the lion used during the fight. To construct what is said_{min} according to the principle of saturation, we start by constructing the literal meaning of the sentence constituents. Since the sentence is ambiguous on how to interpret the possessive relation, the only non-arbitrary way to assign it a single value is to interpret it as what the speaker meant: *“the sword X used in the fight”*. For *“the lion”*, there is a literal interpretation where we interpret it as the dead animal that lies on the ground. Combining these, what is said_{min} by the sentence as a whole becomes something like *“Give me the sword that the lion (the now-dead animal) used in the fight”*. This interpretation is absurd, and clearly does not correspond to any stage in the actual process of understanding the sentence.

This shows that the actual composition of the sentence meaning does not use the minimalist meaning of the constituents. Rather, the composition takes place at the pragmatic level, where contextual information can be used in the composition.

Recanati concludes that it is not possible to construct a four layer model of sentence meaning. Either what is said_{min} is impossible to construct without depending on information from the pragmatic level, or the layer collapses into either the linguistic sentence meaning or the what is said_{prag} level. Additionally, there appears to be no reason why a separate what is said_{min}-layer needs to be part of how we interpret a sentence.

Instead of solving the dilemma between literalism and contextualism, the syncretic approach proposes a model that is harder to defend due to its increased complexity, and in addition inherits the disadvantages of both literalism and contextualism.

2.4 Contextualism

2.4.1 Quasi-Contextualism

The last two positions on the scale are quasi-contextualism and contextualism. Although Recanati positions them in this way, I think they are actually orthogonal positions because they are concerned with different questions. Contextualism (which will be fully introduced in the next sections) is concerned with what is said, and states that there are truth-conditionally relevant contextual influences in what is said that cannot be captured by saturation in the way that the literalist positions want. Quasi-contextualism is concerned with the question of, given that one accepts contextualism, what is the position of the literalist construal of what is said.

The central claim of quasi-contextualism is that the minimal proposition, as proposed by several forms of literalism, could be constructed as a theoretical entity, but that it does not play any kind of role in understanding language or communication in general. Constructing it is therefore nothing more than a futile academic exercise. Alternatives to quasi-contextualism would for example be to claim that a consistent notion of what is said according to literalism does not exist.

Since quasi-contextualism is not really concerned with the main question I am investigating in this thesis, I will not go further into it.

2.4.2 Modulation

As the last position in the spectrum, we get to contextualism itself. This is the position that Recanati takes.

As I have said before, the main thesis of contextualism is twofold. The first part is that no propositional meaning of a sentence can be formed without taking a large—possibly unbounded—amount of contextual information into account. The second is that how the context influences the meaning of an utterance is not something that is controlled by the sentence (i.e. the linguistic meaning), but something driven by the context.

Before we can discuss contextualism as a position, we will have to look deeper into the process whereby this contextual information is incorporated into the sentence meaning. This process is called ‘modulation’.

‘Modulation’ is a term that was coined by other linguists²⁴ to refer to a process

²⁴A. Cruse, *Lexical Semantics*, 1986, pp. 50–3; and Ruhl, *On Monosemy*, pp. 85–95, and

where the meaning of a word in a sentence is dependent on and modulated by the other words in the sentence. For instance “*light*” does not have quite the same meaning in “*a light lunch*” as in “*a light piece of luggage*” or in “*The fire won’t last with such light fuel*”.

Recanati uses the same term not just for how words in a sentence can influence each others meaning, but also for how the larger context can do the same. For Recanati ‘modulation’ is a cover term for all forms of optional primary pragmatic processes. Recanati disagrees with the analysis of “*light*” in the previous example given by Cohen and Ross²⁵ who describes this as the head noun of the phrase influencing the meaning of the adjective ‘light’. According to Recanati and other contextualists, what matters is not what exact words are used, but what the words are describing. Nothing prevents someone from using “*light*” in the same phrase “*light fuel*” to describe not the fuel’s density but its weight.

Another example, given by Searle, shows this form of modulation a bit more extensively: The meaning of the word “*cut*” is modulated differently in the phrases “*John cut the grass*” and “*John cut the cake*”. If I asked John to cut the grass, and he takes out a knife and cuts the grass the way one cuts a cake (or the other way around), he has not done what I requested him to do. However, we can imagine a different situation where asking John to cut the grass I would mean to do it in the way one cuts a cake: Suppose we run a sod farm where we sell strips of grass turf to people to put on their lawn. If I say “*Cut half an acre of grass for this customer*”, I probably do not mean that you should *mow* it, I mean that you should slice it into strips as you could cut a cake or a loaf of bread, so the customer can take it with him.

This shows it is not just the linguistic context of a word that modulates it, the important factor is the *situation* the words are used to talk about. The interaction of word meanings in a sentence, as claimed by Cohen and Ross, is only a reflection of this more general phenomenon.

As the examples show, this modulation is not merely a secondary pragmatic process that adds some truth-conditionally irrelevant ‘shade of meaning’ to what is said, the way literalists would argue. Modulation *is* truth-conditionally relevant. It is the *modulated* sense of words that is used to form what is said.

Another way to interpret these truth-conditionally relevant contextual influences is to say that the different senses of the modulated words used above are not caused by a truth-conditionally relevant modulation, but are a form of disambiguation. Some words have more than one meaning, and determining the sentence meaning requires determining which of a set of possible meanings for ambiguous words is used. This argument does not hold, because just about any

elsewhere. James Ross uses the term ‘meaning differentiation’. Cited in Recanati p. 131.

²⁵On their position regarding modulation, see for example Ross, *Portraying Analogy*; J. Cohen, ‘How is Conceptual Innovation Possible?’, in *Erkenntnis* 25, 1986, pp. 221–38.

word is susceptible to modulation. Often, there is not a fixed set of meanings to choose from, but an unlimited variety. As Dwight Bolinger puts it: “Since the universe never repeats itself exactly, every time we speak we metaphorize.”²⁶

2.4.3 Four Approaches to Contextualism

If one accepts the existence of the process of modulation, there are several views one can have on the issue. Modulation is a process that runs parallel to saturation, being optional and context driven instead of mandatory. The views on modulation are somewhat analogous to the different views on saturation that are expressed by the different positions in the literalism–contextualism debate.

Strong Optionality (SO):

According to this position, modulation is optional. It can happen, but it is not always necessary in order to extract the meaning of an utterance. This view is somewhat similar to the minimalist view on saturation. This view is also the most natural to be held by syncretists or quasi-contextualists.

Pragmatic Composition (PC):

Modulation is optional for individual words, but it is most of the time necessary at the level of full sentences. The meanings of words used in a sentence have to fit together, and normally the unmodulated meanings of two words do not fit. So in practise, modulation happens for nearly every sentence we use.

For instance, in the previous example of “*cut the grass*”, the meaning of “*cut*” has to be modulated to fit whatever is being cut, since lots of things are cut in slightly different ways. Grass, cake, paper, and vegetables are all cut in different ways. Since the unmodulated meaning can match at most one of these forms of cutting (but more likely, is something more abstract that doesn’t match any of these kinds of concrete examples) modulation is necessary in most of these cases. Since the meaning of most words in a sentence need to be adjusted in similar ways, it is fair to say that in practise, sentences in which modulation is not necessary do not occur.

Wrong Format (WF):

Linguistic meaning has a different type or structure from the semantic layer. Modulation is a process that converts one type to the other. Modulation is therefore always necessary because otherwise this conversion does not get made.

Meaning Eliminativism (ME):

This is WF pushed to its extreme. According to ME, there is no linguistic meaning, at least not in anything like the traditional sense. Instead, what

²⁶D. Bolinger, ‘Aspects of Language’, 1968, p. 230. Cited in Recanati p. 131.

words contribute to the sentence meaning is how they were used in past situations. This view is implicitly expressed by Austin and Wittgenstein²⁷.

2.4.4 Meaning Eliminativism

Of the above approaches, we will look at meaning eliminativism especially. It is the most extreme form of contextualism. As an extreme position, one might think that it would not be too hard to find arguments that show its weaknesses, but one reason Recanati looks especially thoroughly at ME is because it turns out to be quite defensible.

ME's main claim, as stated above, is that there is no linguistic meaning in any sense that is comparable to any traditionally known form of meaning. This is best seen if we compare ME to the traditional picture:




Figure 2.1: The traditional contextualist picture

We see how a linguistic meaning is learned by abstraction from encountered uses of an expression. When an utterance is heard in a context, modulation produces the contextualised sense of that utterance, taking the context and the linguistic meaning as inputs.

Meaning eliminativism can be seen as the following scheme:




Figure 2.2: Meaning Eliminativism

The linguistic meaning is entirely cut out, and the processes of abstraction and modulation are merged into a single process. On this view, there is no such thing as linguistic meaning that serves as input to modulation. Instead, the applicability of expressions must be contextually determined, just like the referents of indexicals. Words, qua linguistic types, are not associated with abstract conditions of application, but rather with *particular applications* that were made in the past. This set of past uses is a word's so-called 'semantic potential'.

²⁷According to Recanati. In fact, one could argue that Wittgenstein goes even further, denying the existence of anything resembling a contextualised sentence meaning. That position, however, makes it very difficult to say anything at all about what the meaning of language is, and therefore falls outside of the current discussion.

Some important differences between a word's semantic potential and the traditional view of linguistic meaning are that the semantic potential is not some kind of abstract set of conditions of application, which 'linguistic meaning' is often held to be. Semantic potentials can also include *references* to elements from the actual context of those prior applications, instead of just descriptions of elements.

The idea that truth-conditional content is essentially unstable and that it is therefore impossible to give a definite linguistic meaning for an expression is also implicit in Waismann's and Searle's writings²⁸. For any set of conditions we try to define in order to specify when a certain word or expression applies, we can imagine thinking of a state of affairs in which those conditions hold, and then *embed* that state of affairs into a world that is very different from our own. At some point we need to start asking if, even though our initially specified conditions hold, the words in question are still applicable. For instance, take the well known example of "*The cat is on the mat*". It is not difficult to describe when this state of affairs hold. But now suppose the cat and the mat are freely floating in space, far away from earth. Is it still possible to sensibly speak of the cat being *on* the mat?

Giving an exact definition of the situations in which an expression applies, in such a way that there is no room whatsoever to include additional details in the description, is impossible. I can give a description of my right hand, describing its size, its shape, colour, tissue, the chemical compound of its bones, cells, and so on, but there are always more details that can be added. There is, as Waismann states it, a horizon of description, beyond which more as of yet undescribed details reside. Any properties that are beyond this horizon can be changed using the above embedding trick. It is possible to make a complete description for mathematical concepts, such as the concept of 'check-mate' in chess, but such things are not possible in the real world. Even if we could—theoretically speaking—give an exhaustive description of the position and other state of every atom in the object we are describing, such that according to current physics there is nothing left unspecified, we can not rule out the possibility that in the future the currently accepted theories of physics are replaced by other theories that require additional information that was not in our original description.

We can only specify a limited set of features, while an indefinite number of features remain implicit in the background. Whether a word or expression is applicable depends on how much and in what ways the target situation resembles the 'normal' situation in which the word or expression is usually considered applicable. Going back to the cat on the mat floating in space example, the sentence "*The cat is on the mat*" can be made applicable again by adding some

²⁸F. Waismann, 'Verifiability', in A. Flew (ed.), *Logic and language*, 1st series, Blackwell, 1951, pp. 119–23, cited in Recanati p. 141; J. Searle, 'Literal Meaning', in *Erkenntnis* 13, 1978, p. 211, cited in Recanati p. 142.

more details to the situation:

Suppose we are flying around in a spaceship, and suddenly we see cats and mats floating past us in space. Oddly enough, they are always in cat-mat pairs, and only appear in two attitudes relative to our point of view: sometimes the cat appears to be on the mat, and sometimes the mat is on top of the cat. “Which is it now?” I ask, and you answer “The cat is on the mat”.

Even though in the above example many of the background features are missing that are present in normal situations, like gravity, there are still some features present that make the sentence applicable.

Some Objections

Recanati names two objections that can be made against eliminativist accounts:

Descriptivism objection:

Waismann presupposes descriptivism, i.e. that a definition is a list of descriptions that specify exactly when the expression is applicable. But that is not the only option. We can define “*cat*” as an animal belonging to the same species as *this specimen* (pointing to a nearby cat), or “*gold*” as *this metal*.

Epistemic Limitations objection:

Even though we can not always *know* if a certain word applies to a certain object or not, that is merely an epistemic limitation. Our epistemic limitations do not prevent any word like “*gold*” or “*cat*” from having a definite content and a definite extension.

These objections do not threaten Waismann’s or Recanati’s conclusions regarding the instability of truth- and application conditions. The descriptivism objection would do so if, by using those definitions, we could in fact determine stable conditions of application for words and expressions. But we cannot. Even given the object that is demonstrated as part of the definition, we need to decide if it counts as ‘the same’ as the object we are currently talking about. We need to determine which *dimensions of similarity* are relevant, and to what extent. If we have the piece of metal that was pointed to when defining gold, and some metal-like object we are currently interested in, when do they count as the same? Is the colour important? What about the location, maybe something needs to be in the exact same location as the demonstrated ‘gold’. Does it need to be equally shiny? The same holds for comparing a potential cat to the demonstrated cat. Are castrated cats²⁹ still part of the ‘cat’ species?

²⁹which can clearly no longer produce fertile offspring, the biological definition of being one species

Again, there are an indefinite number of dimensions of similarity that need to be specified, so a definition of a word given in terms of demonstratives is still vulnerable to the same arguments given above that no meaning exists that is remotely like a word's linguistic meaning.

The epistemic limitations objection, if taken seriously, would imply that any nonsense sentence would have a definite meaning. That nobody would know what "*All mimsy were the borogoves*" meant would be no reason to say that it *has* no meaning, its truth conditions would be that it would be true iff all mimsy were the borogoves.

Dimensions of Similarity

To determine if a certain situation warrants the application of a certain expression we need to compare it to other situations in which the expression was used, or, according to traditional views, we need to compare it to the conditions of application that form the linguistic meaning. This linguistic meaning was formed by abstracting out the similarities between different situations in which an expression was used. But, as shown above, determining which dimensions of similarity have to be used is a question that in itself depends on the target context. Another argument to this effect is given by Tversky³⁰:

Whether two items are deemed similar or different depends (among other things) on the 'contrast set', the set of items with which the given items are compared. If we ask which country, Sweden or Hungary, most resembles Austria, the answer depends on the set of countries considered. If that set includes, besides Sweden, Hungary and Austria also Poland, then Sweden will be judged more similar to Hungary, but if the fourth country is not Poland but Norway, Hungary is chosen.

In normal life, the contrast set is usually not explicitly provided but depends on the interests and goals of the conversational participants. It follows that by changing these while leaving the rest of the context the same, one can change the meaning ascribed to a word. For example:

Fred is walking with his young nephew beside a pond where a decoy duck is floating. Pointing to the decoy, he says "That's a duck."

We might ask if this statement is true or false, but we do not have enough information to tell. If Fred has just finished laughing at a sportsman who blasted a decoy out of the pond and he is trying to show his nephew to avoid similar mistakes, then what Fred said is false. But if Fred and his nephew are attending the annual national decoy exhibition, and the boy has been having trouble distinguishing ducks from geese, then what Fred said was true.³¹

³⁰Tversky, 'Features of Similarity', in *Psychological Review* 84, 1997, pp. 327–52, cited in Recanati p. 149.

³¹Travis, *Saying and Understanding*, Blackwell, 1975, p. 51, cited in Recanati p. 149.

Some other arguments

It can be objected that ME does not explain how communication is able to succeed given that no-one has the same linguistic potential for a given word, because everyone has different past experiences. The answer is simple: humans are alike in what they judge as 'similar' or 'different', and thus apply the abstraction/modulation process in similar ways. This allows people to come to similar conclusions regarding what is said.

There are also a few problems that Recanati does not address directly. The first is that our memory is limited, and we are not able to remember every aspect of every past use of a particular word or expression. This makes it psychologically highly dubious to claim that *all* past uses of a word are part of its semantic potential. At the very least some sort of selection or compaction of the sensory input that constituted the past situations must take place. The second problem is that it is impossible to derive a word's applicability, which may depend on an infinite number of dimensions of comparison, from a finite number of past contexts, as is shown by Quine's famous 'gavagai' example³².

Accepting these two problems does not defeat the core claims of ME, but it does require some adjustments. Some sort of selection or filtering process must take place, and there must be some kind of generalisation (but that could take place either when the meaning of a word is learned, or when it is applied). These are not a problem for ME because they do not result in a linguistic meaning that resembles a traditional view, and they still allow references to be part of a word's semantic potential. It does, however narrow the gap between ME and WF, but Recanati already recognises on other grounds that there can be all sorts of in-between positions between ME and WF.

2.5 Recanati's Position

Above, I have given an overview of the spectrum of possible positions on the question of how much what is said by a sentence depends on the context, at least as delineated by Recanati. I will now describe Recanati's own position.

Recanati's point of departure is minimalism, as it is still (together with its close cousin indexicalism) the most widely accepted position. According to him, one of the main problems is that what is said by a sentence according to minimalism has in many cases nothing to do with the meaning we intuitively ascribe to a sentence. The minimal proposition is often not something we are aware of when we use the sentence, and neither does it have to be a part of our intuitively ascribed meaning, as is shown by the examples in section 2.3.3 above.

³²Quine, 'Word and Object', 1960

In order to guard against these kinds of mistakes, both in minimalism and in other positions, we should adopt the principle of availability:

Principle of Availability:

What is said must be intuitively accessible to the conversational participants. (Unless something goes wrong and they cannot be counted as 'normal interpreters'.)

This principle follows first of all from the criticisms on minimalism and the strange kinds of meanings we would have to accept otherwise. A second justification of this principle is based on Grice's idea that saying is itself a variety of non-natural meaning. One of the distinguishing characteristics of non-natural meaning, according to Grice's analysis, is its essential overtness. Non-natural meaning works by openly letting the addressee recognise the speaker's primary intention. One way to do that is by *saying* something, another way is *implying*, which is something distinct.

What is said by a sentence is distinct from what is implied. According to Recanati, this is a clear distinction, but determining what is said can still be a problem. Although what is said is by the above principle consciously available, asking someone what a sentence says may not give us the right answers because unless that person is a linguist who knows about the availability principle, he may not properly distinguish what is said and what is implied. The way to determine what a sentence says, is to present people with situations in which the sentence in question is uttered, and asking them if that use of the sentence was true or false. In Recanati's view, knowing what is said by an utterance directly implies knowing its truth conditions.

Obviously, the details of the situation can have a lot of influence on when an utterance is considered true or false. Therefore what is said is necessarily not a property of a sentence in isolation, but of the utterance of a sentence in a given situation. Contrary to the classical linguistic view that what is said is something that is derived from a sentence (possibly incorporating a limited amount of information from its context by saturation), Recanati adopts a more psychological view in which what is said is the result of a cognitive process, where hearing that something is said is analogous to seeing something. What is said corresponds to the primary truth-evaluable representation that is consciously available to the subject.

Some problems with this notion of 'what is said' are that it is not objective, and that it does not apply to modes of speech like sarcasm, where the speaker does not mean what he says. Recanati's answer to the first challenge is that in determining what is said we should use normal conversational participants, so he excludes participants that make major mistakes or that have some kind of linguistic impairment. Doing so does not make the notion of 'what is said' objective in an absolute sense, but it does make it intersubjective, which is good enough to be usable. Recanati's answer to the second challenge is that in modes

of speech like sarcasm, though the speaker does not mean what he says in a strict sense, for the purposes of this discussion we should include “pretending to say”, or “making as if to say”, as happens in these cases, into the notion of ‘saying’.

2.5.1 Recanati’s proposed model

Recanati also proposes a model of how understanding of an utterance should work. The model is intended to contrast with models such as those proposed by minimalism and other literalist viewpoints. It is still quite sketchy, as Recanati himself admits. He intends for it not to be a fully general model that should be able to answer any question about the meaning of language, but to be a starting point in the discussion and to show that it is at least conceivable that language works the way contextualism and the availability principle say it works.

Recanati’s model, just as in most literalist models, starts out with the literal meanings of the individual words (or other non-decomposable meaning-carrying constituents of a sentence). However, once these are retrieved they are not immediately combined to form the meaning of the expressions they constitute. Instead, the language-interpreting machinery in the brain generates for each word a set of modulated meanings. The potential number of meanings an expression could have is very large, but the larger situational context determines which meanings are accessible enough in that context in order to be considered. Then, when composing the meanings of two words or two expressions into the meaning of a larger composite expression, for each constituent the (possibly modulated) variant of meaning is chosen so that the two constituents fit best, taking the overall situational context into account.³³ This continues up the composition tree of the sentence, until it results in what the sentence says. From what is said, the implicatures are derived.

This model has the desirable properties of being compositional, yet it allows contextual influences to pervasively influence what a sentence says, and it does not require absurd candidate meanings to be derived first before being rejected³⁴.

Let us take a look at one of the previous examples again: “*He wears rabbit*”. The set of candidate meanings for the word “*rabbit*” at first contains the literal meaning of RABBIT STUFF³⁵, but also the stricter meaning of the animal RAB-

³³Recanati here mentions ‘schemata’ as an additional model of when concepts do or do not fit together. See for more information on schemata David Rumelhart, ‘Some Problems with the Notion of Literal Meanings’, in A. Ortony (ed.), *Metaphor and Thought*, (Cambridge University Press, 2nd edn, 1993), pp. 71–82; and ‘Schemata, the Building Blocks of Cognition’, in R. Spiro et. al. (eds.), *Theoretical Issues in Reading Comprehension* (Lawrence Erlbaum, 1980), 33–58. Cited in Recanati p. 36.

³⁴Contrary to Grice, who defends that the derivation of absurd propositions is what causes us to start applying implicatures.

³⁵This is the literal meaning of the mass noun “*rabbit*” according to Recanati. Personally, I disagree, because I think the most salient meaning for the word “*rabbit*” when no specific

BIT, as well as other more derived meanings such as RABBIT MEAT and RABBIT FUR. The word “*rabbit*” could have more distantly related meanings, such as A MEMBER OF THE FOOTBALL TEAM “THE RABBITS”, but such meanings are only considered if they are accessible enough in the given context. When the meaning of “*rabbit*” is composed with “*wears*” (for which of course also a set of candidate meanings has been generated), RABBIT FUR is chosen because that fits best with the accessible meanings of “*wears*”.

There is one problem with this model, and that is that the way it is described here, it is serial. The listener has to first compute what is said, and only then can start to calculate what is implicated, because the secondary pragmatic processes take what is said as an input. There are some reasons to assume this is wrong. First of all, we are not aware of these two things happening serially, we usually become aware of both what is said and what is implicated at the same time. Second, there may be circular dependencies between the implicatures and what is said, because the implicatures could influence which modulated meanings are chosen during the derivation of what is said. Third, implicatures can sometimes be premises as well:

A: “*Why didn't you invite me to the party?*”
 B: “*I only invite nice people.*”

In this example, the fact that A considers B to be not nice is an implicature, but it is also a necessary assumption if B's utterance is to be considered as an answer to A's question.

Recanati's answer is that, while the availability principle holds unmodified for both what is said and what is implicated, for the process of implication that takes place between the two a weaker form of availability is sufficient. Recanati distinguishes three kinds of inferences:

Explicit inference:

This is what is normally understood under ‘inference’. It is conscious and occurrent reasoning which we do by carefully considering each individual step of a deduction in our minds.

Sub-personal implicit inference:

Tacit inference. A system is doing sub-personal implicit inference if it is unconsciously, ‘mechanically’, executing steps which could be described as if it were reasoning, although no real reasoning needs to be going on. This is the kind of reasoning one could ascribe to a computer, but also to a thermostat.

context is specified, is the rabbit animal. That is not a mass term, but I think primary pragmatic processes should allow RABBIT STUFF to be derived from RABBIT.

Personal implicit inference:

This is the same tacit unconscious processing of information as sub-personal implicit inference, except that the cognitive agent to which it is ascribed is *itself* capable of making the inference explicit and of rationally justifying the methods it unconsciously used to arrive at the 'conclusion'. The explicit steps which an unconscious process is 'rationalised' into need not be the exact same steps that the unconscious process originally took. It could be a different deduction that arrives at the same conclusion.

Recanati thinks that the inference of implicatures is of the personal implicit variety. Although the explicit steps need not always be explicitly available, they can be made available by reflection if the listener wants to. The actual processes by which someone arrives at what is said and at the implicatures can be intertwined instead of serial, but upon reflection the listener is capable of making the reasoning that justifies an implicature explicit.

Recanati is aware that his model is still quite sketchy, for example it should be elaborated in order to give a better description of what happens in the unconscious processes in the case of personal implicit inference. However, the model should be able to serve as a starting point for further investigation, and it shows that an alternative to the more mainstream approaches of literalism is at least conceivable.

Chapter 3

Critical Evaluation

In this chapter, I will evaluate Recanati's book and the position he takes.

Generally speaking, I agree with Recanati and his criticism of literalism. I agree that there are lots of constructions in language of which the meaning depends on the context in which it is used, and just applying saturation does not result in a form of meaning that agrees with our intuitive notion of what sentences say. Neither does this minimal proposition appear to play a role of any importance in the process of analysing a sentence meaning. Several examples show that context comes into play when the meanings of sentence constituents are composed, so I do not think any approach will be successful in which the context is only considered after the composition of lexical meanings is completed.

There are four general points of criticism I have with Recanati's book and contextualism in general. The first is the book's lack of structure. But since I have already discussed this point I will not go into it again. My second point is that there are several developments in literalistic approaches that Recanati does not mention. The third is that I think Recanati's arguments sometimes apply only against specific forms of a theory, while he could make more general arguments that are applicable against a broader class of ideas. My last point of criticism is not one against Recanati specifically, but against the position of contextualism in general. While literalist approaches have resulted in lots of mathematically formal theories (way too many, one might say), the theories from the contextualist side of the spectrum remain sketchy and informal. While contextualism has some excellent points regarding what is wrong with literalism, when it comes to testable mathematics they do not offer viable alternatives.

In the remaining of this chapter I will discuss these four points, and also offer some future research ideas regarding the last point.

3.1 Other Approaches in Minimalism

Recanati does his best to answer a lot of objections that have been made against him and other contextualists. However, there are still some major developments he does not give any attention to. Recanati is not entirely to blame for this, since discussing all existing approaches and answering all arguments that have been made can not be done within the size limitations of a single book.

In the examples Recanati gives when discussing minimalism, he assumes a rather straightforward way of how for example quantifiers, definite descriptions and numerals work. In the example of *“The table is covered with books”* Recanati only discusses a Russellian analysis of the definite article *“the”*, even though this analysis is not very widely used anymore. An alternative analysis is to distinguish the statements made by a sentence from the preconditions it has, and requiring that if a sentence is to be applied correctly, its preconditions must be satisfied. In such an analysis, the sentence *“The table is covered with books”* would have a precondition that a certain table is already identified in the discourse or general context. Whether such an approach should be considered as part of the literalist or contextualist side of the spectrum depends on the details of the theory, but theories that are clearly literalistic have been constructed. One can imagine that both a sentence’s minimal proposition and its preconditions are generated in a purely bottom-up fashion, and then checking of the preconditions is done.

There has also been further research on quantifiers since the time of Russell. A general framework for quantifiers should be able to also handle non-absolute quantifiers like *“many”* and *“most”*, rather than only *“some”* and *“all”* as standard first order predicate logic does. Frameworks that deal with non-absolute quantifiers probably also handle absolute quantifiers differently, which potentially invalidates Recanati’s argument. As I already mentioned³⁶ Recanati makes the assumption that a quantifier by default quantifies over everything, an idea that originated in logic but that does not necessarily have to be taken for granted in natural language analysis.

Some other approaches that Recanati does not discuss are Montague Grammar and Discourse Representation Theory. I do not claim that Recanati’s arguments cannot be applied against them, but doing so is not always a simple straightforward thing. Recanati’s arguments assume that what a sentence says is something truth-conditional, and therefore is to be modelled as a formula of some sort of logic system. But in Montague Grammar a meaning is not a logical formula, but a set of possible worlds. While it is probably possible to translate Recanati’s arguments into a form that is usable in possible world semantics, doing so is not entirely trivial.

Discourse Representation Theory (DRT) is an approach whose main aim is to model anaphora resolution (including descriptions using the definite arti-

³⁶in footnote 16 on page 14 regarding minimalism

cle “*the*”) by building up a mathematical representation of a discourse. This (representation of) discourse then determines which referents are available to resolve an anaphor. In some aspects, this is a very contextualistic approach. DRT certainly tries to take the context into account. On the other hand the context that DRT deals with is only the discourse as represented in the text that a specific sentence is a part of, and not the general, wide context. But doing that is currently beyond the practical possibilities of mathematics. DRT does not fit cleanly into one of the categories that Recanati delineates. It deals with a limited amount of context, but what it does is more than just a straightforward filling in of open variables. As far as I am aware, DRT does not use the context when composing a sentence meaning out of the meaning of constituents.³⁷ This places it on the literalist part of the spectrum, but with a lot of relatively complicated mandatory pragmatic processes. I would say that a form of indexicalism would be the best fit, but DRT proponents may argue differently, especially more theoretically focused ones for whom computational practicality is not so much of a concern.

3.2 More General Lines of Argumentation

Recanati offers a number of arguments against different forms of literalism. He often targets very specific forms of literalism, but I think he could in some cases give arguments that are more broadly applicable.

In the case of indexicalism, Recanati specifically targets the two discussed criteria for deciding when a contextual influence is saturation and when it is not, without actually arguing against indexicalism in general. A natural reaction from an indexicalist (assuming he accepts Recanati’s arguments) would be to go looking for a better criterion.

One step up, Recanati splits his argumentation against literalism at large into several arguments against the different forms of literalism. Recanati actually also offers arguments that are valid against literalism in general, or at least against a broad group of literalist positions, but they are all discussed in the context of specific forms of literalism, which makes it not obvious that they are more widely applicable.

Against indexicalism, Recanati could make a more general claim: no matter what criterion is used to distinguish saturation and non-saturation contextual influences, either some influences will be classified as pragmatic while in fact influencing the sentence’s truth conditions, or the theory will result in an unbound number of variables in the formula that represents what is said. The

³⁷See for example <http://svn.ask.it.usyd.edu.au/trac/candc/wiki/boxer> (retrieved 05-07-2011). This is an implementation of DRT that takes as input a parse tree that is generated by other tools, so its discourse representation does not influence parsing. However it could be argued that that is merely a practical limitation of this system, which could be lifted in the future.

latter implies that nearly every contextual influence is handled by saturation, including ones where everyone intuitively agrees they are pragmatic.

I think such a claim is true, but it is difficult to defend this since it is such a general claim. But I do not think it is impossible. One could argue as follows: Waismann and Searle claim that one can embed any 'normal' situation in an abnormal surrounding or world so that the truth of any claims about it would be questionable.³⁸ This is done by changing some of the properties of the new world that are normally taken for granted. There are an unbound number of dimensions of similarity, and for any sentence that says anything about the physical world it is very hard—if not impossible—to establish an upper bound on how many of these dimensions can influence its truth conditions. Since each such dimension requires a variable in the indexicalist representation of what is said, the number of variables is also unbounded.

When discussing syncretism, Recanati makes an argument to show that the context is needed not just after the lexical meanings of words have been composed together, but also during the composition. This argument concerning the lion's sword (see subsection 2.3.5: 'Other Problems' on page 22), is applicable against other conceptions of the minimal proposition as well. Recanati could have made this more obvious, and in that way make a firmer argument against literalism in general.

3.3 Formal Rigour and Future Research

Despite some shortcomings, I find the argumentation put forth by Recanati and other linguists and philosophers of language with a contextualist signature very convincing. Many of the examples show that context plays a much larger role than many less contextualistic theories will recognise. I often like to think that other philosophers and linguists, when presented with these kinds of arguments, would also be convinced by them. Why, then, is literalism, specifically forms somewhere between minimalism and indexicalism, still the theory of choice for the mainstream?

The one thing in which I think contextualism is severely lacking, is formal models. The literalist tradition has brought forth a plethora of mathematical formalisations of their ideas, for example X-bar theory, Chomsky's minimalism theory, Montague grammar, and DRT. While these models usually fail when presented with the kind of examples Recanati demonstrates, at least literalists *have* models with which they can analyse a sentence or which they can implement in a computer. Formalisation under contextualists is very sparse in comparison. Their ideas often do not transcend the stage of philosophical argumentation.

Recanati's book is no exception. In some places Recanati tries to give a start to a formalisation, for instance when he discusses the *Circ* operator, but these formalisations are not anywhere near the level needed to actually be able

³⁸See subsection 2.4.4: 'Dimensions of Similarity' on page 29.

to apply and check them. This is recognised by Recanati, saying that he first of all wants to show that such analyses are at least conceivable

The position of meaning eliminativism, as described by Recanati, distances itself even further from the possibility of formalisation. By denying the existence of some kind of linguistic meaning, they force a language philosopher to deal with the highly complicated history of uses of an expression, instead of a more simple abstraction from those.

There are reasons for this lack of formal models, of course. The main one being that a formalisation of a contextualist theory will usually require a formalisation of the general context in which language is used. Since this wider context is not limited in what is part of it, this amounts to formalising the entire world in as far as humans understand it, including human thought. At the current level of scientific capability this is far too much to ask for. Doing so probably equates to building a general artificial intelligence.

Since that is not possible, linguists and language philosophers do what scientists usually do when confronted with problems that are too complex to handle: they simplify. One popular definition of a 'model' is 'a simplified representation of a part of reality'. If literalist models are seen in this light, instead of as fully accurate representations of what happens in reality, they make a lot more sense. From this point of view literalism is a necessary simplification of how language works, that is needed to keep it manageable.

So, does this mean that contextualists should resign their positions? Not necessarily. By their lack of putting forth viable formal models, contextualists often take up a negative position (or at least a not-very-positive position). This in itself has value since it is important to be aware of the limitations of literalistic models. But if our understanding of human language is to move forward, positive positions are necessary as well. Formalising general wide context is not possible, but that does not have to stop contextualism. Contextualists could try to come up with limited models, that only try to formalise some very small amount of context. Such a model will not be able to give a good analysis of arbitrary sentences, only of sentences that are about that piece of context. But, as long as these small formalised contexts are not seen as islands but as things that are meant to be connected and grow, our understanding of language in context can grow as well. Even if enlarging the amount of formalised contextual information is difficult, these models can help in understanding how the process of language understanding works, such as what computation and information flow occurs, where feedback is necessary, and so on. I think this is something that deserves attention from future research.

Index

- binding criterion, 17, 18
- Binding Fallacy, 18
- binding fallacy, 19
- bottom-up pragmatic processes, 9
- Cohen, 24
- contextualised meaning, 6
- contextualised sense, 6, 26
- Contextualism, 4–7, 19, 23, 25, 26, 32, 35, 38, 39
- contextualism, 23
- contrast set, 29
- definite description, 8, 13, 36, 37
- demonstrative, 5, 8, 29
- dimensions of similarity, 28, 29, 38
- Discourse Representation Theory, 36–38
- DRT, *see* Discourse Representation Theory
- eternal sentence, 8, 10, 11, 15
- explicit inference, 33
- first order logic, 7, 8, 10
- formal language, 7, 8
- free enrichment, 11, 12
- horizon of description, 27
- Ideal Language Philosophy, 4, 7, 8
- implicature, 6, 9, 12, 13, 32–34
- indexical, 5, 8, 9, 16, 17, 20, 21, 26, 37, 38
- Indexicalism, 7, 10, 14–19, 30, 37, 38
- linguistic meaning, 5, 6, 9–11, 15–17, 20, 23, 25–27, 29, 30, 39
- Literalism, 4–7, 19, 20, 23–25, 32, 34–39
- loosening, 11, 12, 20
- mandatory pragmatic processes, 9, 10, 22, 37
- ME, *see* Meaning Eliminativism
- Meaning Eliminativism, 25–30, 39
- minimal proposition, 10–15, 19, 20, 23, 30, 35, 36, 38
- Minimalism, 7–16, 19–21, 30–32, 36, 38
- modulation, 23–26, 30
- Montague Grammar, 10, 36, 38
- optional pragmatic processes, 9, 13, 15, 16, 22
- optionality criterion, 15–17
- Ordinary Language Philosophy, 4, 7, 8
- PC, *see* Pragmatic Composition
- personal implicit inference, 34
- Pragmatic Composition, 25
- pragmatic process, 8–12, 21
- pragmatics, 4, 19
- pre-propositional, 9, 10
- presupposition, 8, 18
- primary pragmatic processes, 9, 10, 19, 24, 33
- Principle of Availability, 31
- proposition, 5, 9, 10, 13–15, 17, 21, 23, 32
- quantification, 8, 14
- Quasi-Contextualism, 7, 23
- Ross, 24

- Russell, 4, 7, 8, 13, 36
- saturation, 9–11, 13–15, 17–23, 25, 31, 35, 37, 38
- Searle, 24, 27, 38
- secondary pragmatic processes, 10, 19, 21, 24, 33
- semantic potential, 26, 27, 30
- semantic transfer, 11, 12, 20
- semantics, 4, 16, 19, 23, 36
- sentence meaning, 5, 6, 19, 21–24, 26, 35, 37
- SO, *see* Strong Optionality
- speaker's meaning, 5, 6, 13
- specification, 12
- speech act, 4
- Stanley, 17, 18
- strengthening, 12
- Strong Optionality, 25
- sub-personal implicit inference, 33, 34
- Syncretism, 7, 19–23, 38
- top-down pragmatic processes, 9
- Traditional Literalism, 7, 8, 10
- truth condition, 4, 7, 9, 11, 15, 20, 29, 31, 37, 38
- Tversky, 29
- Waismann, 27, 28, 38
- WF, *see* Wrong Format
- what is implicated, 6, 33
- what is said, 5, 6, 8–11, 13, 14, 19–24, 30–34, 37, 38
- Wrong Format, 25, 30