Truth-Theoretic Contextualism: Dissolving the Minimalism/Contextualism Debate

MSc Thesis (*Afstudeerscriptie*)

written by

Wouter Beek
(born June 6st, 1983 in Swalmen, Netherlands)

under the supervision of **Prof. Dr. Martin Stokhof**, and submitted to the Board of Examiners in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam*.

Date of the public defense:	**Members of the Thesis Committee:**
August 24, 2009	Prof. Dr. Martin Stokhof (supervisor)
	Dr. Paul Dekker (second corrector)
	Prof. Dr. Ir. Remko Scha
	Prof. Dr. Veltman (Board of Examiners)

Institute for Logic, Language and Computation

TABLE OF CONTENTS

Page

Chapters:

CHAPTER 1

INTRODUCTION

The scope of the research field of semantics seems to shrink in the recent debate between the viewpoints of contextualism and (semantic) minimalism. Minimalists hold that semantics is the formal endeavour that derives a minimal proposition from the syntactically structured input string. This semantic content is then fed to the subsequent pragmatic module in which meaning is supplemented in a non-formalisable and non-tractable way. Since most of what we intuitively think of as the content that gets meaningfully communicated in making situated utterances belongs to the output of the latter, pragmatic module (i.e. the so-called 'speech act content'), minimalists argue that semantics is only minimally relevant to the pursuit of formulating a theory that explains successful communication.

Contextualists maintain that in the minimalist endeavour, the realm of the semantic (i.e. the formal approach to meaning) has been narrowed down to such an extent, that what is left is not even interesting with respect to the elucidation of communication. They therefore conclude that the formal, semantic approach to meaning – at least from the point of view of communication – should be done away with in theorizing. What replaces the formal tradition is a plurality of casually formulated principles that purport to describe acts of communication in abstract terms, e.g. 'that the communicated content be relevant to the addressee'. These principles, although interesting in themselves, are neither mutually interrelated, nor stringently verified with respect to empirical instances of communication that might invalidate them.

In this essay I will be concerned with the minimalism/contextualism-debate from a methodological point of view. Although the here espoused methodological considerations will have a profound influences upon the descriptive linguistic practice, the discussion will proceed along conceptual lines.

As the minimalism/contextualism-debate has progressed, some minimalists have joined the contextualists in the opinion that formal semantics cannot explain communication (not even to a meagre extent). For instance the small bit that Borg [6] takes semantics to consist of, after all the critique of the contextualists has been taken to heart, is unlikely to make many philosophers of language very happy. She says that her book "turns out to be a denial of the bunch of views often labelled 'the linguistic turn'."[1] And all the while we were thinking that, finally, philosophy would get somewhere!

As is apparent from the above quote by Borg, in the discussion over semantic content there seems to be a tendency, on both the minimalist and the contextualist side, to close down most of the semanticist's workspace in philosophy. But there are still so many tools in that workspace none of the crude devices of any of the other philosophical denominations can match. Or so I will argue.

[1] See [6, p. 12].

I think that abolishing the traditional endeavour of formal semantics with respect to the elucidation of communication, from a theoretical point of view, is premature. But what is more, a theory of communication can only be legitimately called *scientific* provided the analysis of meaning that drives it conforms to formalistic strictures (such as those laid down by Tarski, see Section 2.2). Given the fact that there is a potential infinitude of distinct utterances of which communication might consist, only a formal mechanism will be able to guarantee that we are able to explain the meaning of *every* natural language sentence. This is why, in keeping in line with the formal tradition, I maintain that a theory of the transmission of meaning (i.e. communication) must (amongst other things) adhere to the *formalization criterion* (which is explained in Section 2.2).

But I will not merely be defending tradition here. I will make fundamental alterations in the light of worthwhile considerations that can be drawn from the minimalism/contextualism-debate. This will result in a reconceptualization of the traditional approach towards semantics, one in which language is viewed upon as inherently intertwined with the human purpose of information interchange. The way in which a theory of meaning and a theory of communication ought to be interrelated was poignantly formulated by Pagin and Pelletier [35, p. 35] in the following way:

> A natural methodology for justifying a semantic theory S of a natural language L is to see S as part of a more comprehensive theory C of communication by means of L.

The purpose of this essay is to live up to this programmatic formulation.

The reason why I do not want to treat of semantics as exempt from communicative use, besides its being arguably foreign to the essence of language, is that such an account cannot be linked to methods of empirical verification. Once we observe the dynamics of linguistic interchange, it is clear that whatever semantic theory one might come up with, it must be able to explain how meaning changes in the face of a change in language use. Moreover, while formal strictures underdetermine the semantic effort, and while intuitions as to what individual sentences mean do not provide a steady beacon, a semantic theory that does not take into account empirical instances of communication, has – scientifically speaking – no objective verification criterion to go by. And yet it is crucial to check the theoretical claims we are making against concrete instances of what we purport to be explaining in the first place (i.e. specific instances of communication).

A further distinction we need to make is that between what Davidson calls 'the uncovering of logical structure' and 'the analysis of concepts' (see Section 2.3.1). The former is concerned with positing just enough structure in order to form a truth-based theory of meaning. It is this notion of semantics that we shall be concerned with in the present essay, since it effectively narrows down to that part of semantic theorizing that suffices for the explication of communication. The latter kind of semantics, i.e. the one termed 'the analysis of concepts', digs much deeper and tries to unravel the logical structure of words and phrases far beyond the point that is needed for explaining

communication.[2] [3] From the foregoing it is clear that in this essay I will only be concerned with the former strand of semantics.[4]

I will take it to be a defining characteristic of communication that whatever concepts the theory puts forward in order to explain an interpreter's understanding of a speaker's speech, the ingredients the interpreter must attain in order to be able to do this, must be readily available to that interpreter. This means that a theory must adhere to what I call the *principle of communication*:

Principle of communication: A speaker can only communicate what an interpreter may be likely to catch up on.

I will adhere to this principle constantly. Most notably it is an inherent property of the methodology of radical interpretation that I will make extensive use of. Because this methodology echoes the same assumptions that I have put forward above, I will formulate my solution to the minimalism-/contextualism-debate in terms of truth-theoretic semantics and radical interpretation. These will be extensively discussed in Chapter 2. In Chapter 3 I will provide an overview of the recent minimalism-/contextualism-debate. In Chapter 4 I will combine the resources of the foregoing two chapters in redefining truth-theoretic semantics, so that it will incorporate the contextual supplementation of meaning. Finally, in Chapter 5 I will evaluate the ability of the various theoretical frameworks, that have been introduced in the preceding chapters, to explain communication.

[2]The branch of semantics that is termed 'the analysis of concepts' can have many purposes. For example, if someone is interested in epistemological matters, one might undertake an analysis of the verbs 'know' and 'believe' (and their interrelation). For purposes of an ethical nature, one might look into the fine-structure of words like 'good' or 'must'.

[3]Cappelen and Lepore, who see it as essential to a viable account of (minimal) semantics that it explains communication, draw a distinction between semantics proper, i.e. what is needed to uncover the minimal proposition, and what they – somewhat misleadingly – call 'metaphysics' (with which they characterize considerations that do not legitimately belong to the field of semantics).

Although their distinction seems to mirror the above one, it must be remarked that on my reconceptualization of semantics the distinction is itself conditional on the language for which meaning is given. This implies that it is only relative to the practices the purpose of which language serves that we can dissect what is necessary for communication from what is necessary with respect to the elucidation of others aspect of human action (i.e. what Cappelen and Lepore call 'metaphysics').

[4]The distinction that is drawn here may seem evident, but it is not always respected. An instance of this is Borg's claim that her approach to semantics invariably denies all of what 'the linguistic turn' has brought us (see the quote above). Even if all she claims is right (which it is not the case, of course), then this will still only affect that part of the linguistic turn that is concerned with the notion of communication.

CHAPTER 2

TRUTH-THEORETIC SEMANTICS

2.1 Overview

In this chapter we will draw the contours of a truth-theoretic semantics. The explication will proceed along the lines laid down by its inceptor Donald Davidson, since it is in his original formulations of the matter that the topic of communication, which is pivotal to the present endeavour, is borne out most clearly:

> [...] a theory of truth shows how we can go from truth to something like meaning – enough like meaning so that if someone had a theory for a language verified in the way I propose, he would be able to use that language in communication. [17, p. 74]

At the centre of a truth-theoretic semantics stands a theory of truth. Tarski [52] showed how to construct such a theory of truth for the formal language of (what we would now call) first-order predicate logic. In order to use such a truth theory for interpretation, it ought to be adapted so as to be applicable to natural language expressions. Traversing from a theory that defines truth for some formal language to a theory that defines truth for (the sentential expressions of) some natural language, is therefore the main hurdle in constructing a truth-theoretic semantics.

In Section 2.2 we shall first describe Tarski's theory regarding the specification of truth predicates for formal languages. In Section 2.3 we show the ways in which Davidson sought to alter Tarski's work so as to make it possible to define truth predicates for natural languages. In Section 2.4 we discuss the thought experiment of radical interpretation, in which we explain how a truth theory is to be evaluated in the light of empirical evidence, in order to assure that the truth theory indeed gives the meaning of the sentences for which it is defined. In Section 2.5 we provide an overview of the way in which Davidson sought to incorporate contextual factors that impinge on meaning into the truth definition. This last section will thus introduce the topic of this essay, i.e. the contextual dependency of meaning, which will be further elaborated upon in subsequent chapters.

2.2 Truth theories for formal languages

A Tarski-style truth theory for a formal language consists of a set of axioms Σ that puts enough restrictions on a predicate 'is-T-in-F', thus constraining that predicate's extension to all and only the true sentences occurring in some formal language F.[1] [2]

The criterion of correctness, that guarantees that what has been defined is in fact truth for F, is that for all and only the sentential expressions of F, we can deduce from Σ as a logical consequence a corresponding *T-sentence*. The schema for such a T-sentence is (2.1).

$$s \text{ is-}T\text{-in-}F \text{ iff } p. \tag{2.1}$$

This schema is instantiated to a specific T-sentence by replacing the s with a structural description[3] of an object language sentence (this is the language the theory is meant to define truth conditions for, in this case the formal language F), and by replacing p with a meta-linguistic sentence that correlates with s (where the meta-language is the language in which the theory is couched, in this case English plus some additional logical symbolism).[4] The meta-language correlate that is filled in for p is either a sentence translating the object language sentence referred to by the name filled in for s, or it is that named sentence.[5]

That a truth definition has all and only T-sentences for the formal language we are defining truth for (i.e. the language named by what fills in for F) as its logical consequences is called the criterion of *material adequacy*. It makes sure that the truth definition is in line with the intuitive grasp of truth (since it is intuition that is the arbiter in establishing whether a T-sentence that is proven by Σ is a legitimate consequence or not).

The requirement that a theory of truth must entail, for each and every sentential expression in the formal language F, a T-sentence like (2.1) is called *convention T*. (A set of axioms Σ for some formal language F is said to adhere to Convention T, and is said to satisfy the material adequacy criterion, iff it is a truth theory that entails all and only true T-sentences for the sentences that occur in F.)

[1] With a language, whether it be a natural or a formal one, we mean a set of sentences (which can be either finite or infinite).

[2] It is a characteristic of Tarski's semantic notion of truth that it can only be defined for some specific language. This means that there is no general notion of truth, i.e. one that transcends all languages. What does transcend all language is the intuitive grasp of truth that – Tarski argued – Aristotle had in mind (see [52, p. 155]). The only thing that is common to all truth predicates is that they all answer the material adequacy criterion.

[3] A *structural description* of a sentence is a concatenation of its sentential primitives (i.e. the individual signs from which the sentence is composed).
Since the number of sentences in F may be (and often is) infinite, we need a procedure for generating such descriptions from a finite set of primitives and a finite number of structural transformation rules that operate on both primitives and complexes (i.e. a recursive specification of syntax).

[4] In order to establish a correct truth theory it is not always sufficient to provide a recursive definition for F (although this is of course necessary). For not only do we (i) need to construct each and every expression in F by finite means, but we also (ii) need to pair each of those expressions with their respective truth conditions (as formulated in the meta-language) by equally finite means. Now for a Tarski-style truth theory, requirement (ii) is automatically met, since not only the syntax but also the semantics can be recursively defined. Moreover, there is a direct link between every step in the recursive construction of the syntactic expressions filling in for s and the recursive construction of its semantic correlates filling in for p (this is the way in which most logics are defined).

[5] In order for this latter condition to hold, the object language must be a proper subset of the meta-language.

But the criterion of material adequacy – in itself – does not guarantee that the definition provides any more clarity than the original intuition of truth does. A further criterion is therefore needed in order to guarantee that the truth definition is not only *correct* but also *theoretically informative*.

Tarski's second criterion, which is to guarantee theoretical informativeness of a truth theory, is that the language in which the concept of truth is to be defined (i.e. the meta-language) must be *formally specified*.

A language is formally specified (or formalized) if its structure is specified solely by an appeal to the *formal aspects* of the expressions that comprise the language. The specification of the formal aspects of the expressions of a language consists of (i) an enumeration of all the primitive terms of the language (i.e. mostly words), and (ii) composition rules that allow defined terms to be construed out of other defined terms and/or primitive terms. These two components allow the automatic construction of all the expressions in the language. Other properties that are necessary in order for the formalization criterion to be met are (iii) a method for distinguishing, from among the thus generated set of expressions, those that are sentential in nature, and (iv) a specification of the deductive structure of the language, i.e. a description of the set of theorems (or provable sentences) in the language.

As a final condition of formality, (v) a truth theory's T-sentences are not allowed to contain words denoting semantic concepts the theory ought to elucidate in the first place (e.g. the word 'satisfies' – introduced in Section 2.2.2 – is not allowed to occur anywhere in the T-sentence[6]).

All that a truth theory accomplishes in terms of making the notion of truth more explicit, more clear, and finitely introspective, all derive from this formalization criterion, while the criterion of material adequacy makes sure that in the process of formalization the original and intuitive notion of truth is preserved.

2.2.1 The semantic paradoxes

There is still one aspect of the formalization criterion that has as of yet been left unspecified. Since our theory involves both sentences and names referring to those sentences, it becomes – unless we pose further restrictions – possible to formulate so-called *semantic paradoxes*.

The problem arises whenever the meta-language is *semantically closed*, i.e. whenever it contains (i) the expressions for which truth is defined (i.e. the sentences of object language F), (ii) names for the expressions for which truth is defined (i.e. names for the sentences in F; the things that are to fill in for the s in the T-sentence scheme), (iii) semantic terms like 'true' and 'satisfies', (iv) the assurance that all the sentences that determine the extension of the predicate 'is-T-in-F' are present in the meta-language (i.e. the condition that the meta-language is able to translate every sentential expression from the object language), and (v) some marginal logical terminology, most notably the 'iff'.

[6]It is not allowed in the object language sentence because of semantic closure (see Section 2.2.1). Also, it is not allowed in the meta-language sentence because it will not translate a concept from the object language. Finally, it is not allowed in the rest of the T-sentence because this would presuppose what ought to be explained.

Tarski makes sure the problem of semantic closure does not occur within a truth definition by locating element (i) in the object language, and the elements (ii)-(iv) in the meta-language (element (v) is allowed to occur in both[7]).

Most strikingly, the thing we are defining, i.e. truth, can itself not be represented by a predicate in the language we are defining it for.[8] Davidson expresses this in saying that "there may in the nature of the case always be something we grasp in understanding the language of another (the concept of truth) that we cannot communicate to him."[9] (The same considerations apply for the satisfaction predicate, see Section 2.2.2.)

2.2.2 Satisfaction

Because of the recursive specification of the object language, which is recursively in line with the assignment of truth conditions in the meta-language, the notion of reference is necessarily invoked. Let us elaborate on how exactly this dependency is instituted.

If the truth definition for the sentential expressions of the object language is to proceed by finite means, then it will not suffice to (i) provide truth definitions for all the atomic expressions, and (ii) specify some finite mechanism for constructing truth definitions for complex sentences on the basis of those for atomic and (other) complex sentences. This is evident from the fact that in general complex sentences will not reduce to *sentences* of a lower complexity (take for instance the ordinary construction rules for quantified sentences, where the direct parts are variables and open sentences[10]).

Another complicating factor is that there need not be a finite number of atomic sentences in the language. Even under the assumption that there is a limited number of predicate and constant

[7]It is apparent from T-sentence (2.2) that the identity of the meta-linguistic translation of the object language connective and the connective that is used in the definition, is not in opposition with providing the right truth conditions for the object language sentence containing the connective 'iff'.

$$\text{`}p \text{ iff } q\text{' is-}T\text{-in-}L \text{ iff } p \text{ iff } q. \tag{2.2}$$

[8]Why can we not include (iii) into the object language? Suppose the truth predicate 'is-T-in-F' would occur in the object language, then the it could contain a sentence like (2.3).

$$\text{The Liar is not } T\text{-in-}F. \tag{2.3}$$

We give (2.3) the names 'The Liar' and 'The Liar is not T-in-F' (the latter is the structural description that the meta-language uses standardly in order to refer to sentence (2.3) in the object language). Since the meta-language contains both (ii) and (iv), we can substitute (2.3) for 'The Liar' in (2.4), resulting in (2.5), which leads to a contradiction.

$$\text{`The Liar is not } T\text{-in-}F\text{' is-}T\text{-in-}F \text{ iff The Liar is not } T\text{-in-}F. \tag{2.4}$$

$$\text{`The Liar is not } T\text{-in-}F\text{' is-}T\text{-in-}F \text{ iff `The Liar is not } T\text{-in-}F\text{' is not } T\text{-in-}F. \tag{2.5}$$

The consequence of this argument is that the object language can never contain a term translating the meta-language predicate 'is-T-in-F'.

[9]Davidson [14, p. 29].

[10]Provided one sees the substitution interpretation of quantified formulas as a viable alternative, complex sentences – as far as they make use of quantificational structure – can be made to reduce to atomic sentences exclusively. We do not pursue this strategy though, because the substitutional account violates requirements that are central to truth-theoretic semantics as defined by Davidson. It would however be of interest to investigate how much of the original theory could be left intact under the adoption of the substitution paradigm, but this is a study in itself and will thus not be pursued here.

symbols in F, the number of atomic sentences (consisting of an n-ary predicate symbol and n term symbols as arguments) may be infinite. The reason for this is that complex predicates can be build using logical connectives and variables. Also, complex singular terms can contain arbitrary nestings of functional operators. Therefore, in order to conform to the requirement of the truth theory's finitude, a recursive specification of the truth conditions of sentences will have to appeal to something that is more basic than sentences.

Tarski [52, p. 189] took the more basic (i.e. sub-sentential) elements to be sentential functions and tuples of objects that satisfy those functions. Formation rules, governed by the semantic principle of satisfaction, allow the (potentially infinite) total set of atomic sentences to be construed.[11] Adding to this construction rules for the formation of complex sentences out of atomic ones, makes one end up with a finite specification of an infinite number of sentences, including the seemingly problematic instances of quantificational structure and complex terms.

2.3 Truth theories for natural languages

Tarski intended his definition of a truth theory, described in Section 2.2, to be used for formal languages exclusively. He nevertheless intimated that the same kind of truth definition might be extended to apply to natural languages also (albeit only approximately):

> *The problem of the definition of truth obtains a precise meaning and can be solved in a rigorous way only for those languages whose structure has been exactly specified.* For other languages – thus, for all natural, 'spoken' languages – the meaning of the problem is more or less vague, and its solution can have only an approximate character. Roughly speaking, the approximation consists in replacing a natural language (or a portion of it in which we are interested) by one whose structure is exactly specified, and which diverges from the given language 'as little as possible'. [51, Section I.6]

The only requirement that, according to Tarski, ought to be instituted in order to make the theory deal with natural languages instead of formal ones, is to make it conform to the formalization criterion (as explained in Section 2.2). The formally specified language, Tarski suggests, should *replace* the natural language. However, as long as the formal language replacing and the natural language replaced can be related to one another in a systematic way, linguistic analysis without language reform is possible.

In line with formal languages, a truth theory Σ for a natural language L must consist of a formally described set of axioms from which T-sentences for each and every sentence in L logically derive. In order to make a finite stock of means cover an infinite stock of ends, as we did earlier for formal languages, we need to relate an infinite number of mentioned expressions that fill in for s with an infinite number of interpreted expressions that fill in for p in the T-scheme.[12]

Since the meta-language must be rich enough to state the truth conditions for all of the object language expressions, it would seem natural to let it just *be* the object language (plus the truth

[11] Atomic sentences that contain only constants do not need the principle of satisfaction, and are construed in the original way.

[12] The reason why we need a *mechanism* that generates *all* the sentences of a language (probably recursively), is that the names that refer to sentences must be construed in a systematic way. We simply cannot make names for all the sentences that fill in for s in the expression 's is a true sentence in L' in an ad hoc fashion.

and satisfaction predicates).[13] A recursive specification of the object language sentences, that is used in giving the names that fill in for s, could then also be used to indicate the way in which the meta-language sentences that fill in for p can be structured so as to make the specification of truth conditions parallel the composition of both object language names and meta-language sentences.[14]

2.3.1 Logical structure and the analysis of concepts

An important distinction that Davidson repeatedly invokes in his specification of a truth-theoretic semantics, is that between uncovering *logical structure* and giving an *analysis of individual concepts*.[15]

> I think it is hard to exaggerate the advantages to philosophy of language of bearing in mind this distinction between questions of logical form or grammar, and the analysis of individual concepts. [14, p. 31]

The larger part of the research tradition called 'philosophy of language' is concerned with extracting philosophical implications from the analysis of the way in which words are used in linguistic intercourse. For instance, the way in which the word 'good' is used is thought to reveal something about ethics; another example is the supposition that if only we knew the legitimate assertion conditions for the words 'believe' and 'know', we would have made headway in construing a theory of epistemology.

But Davidsonian truth-theoretic semantics is not concerned with such strands of conceptual elucidation. It is not part of the bargain of explaining meaning that one also extract, as a by-product say, philosophical theories of what it is to be good or descriptions of the conditions for genuine knowledge attributions. The only philosophically interesting concept a theory of meaning ought to elicit is that of *communication*.[16]

The criterion for grouping some aspect of grammar under the header of either logical structure or conceptual analysis, is not to be defined independently of the notion of successful communication. An aspect of grammar can be taken to belong to logical structure iff it is necessary for the truth theory that functions as the meaning theory for the natural language under consideration.[17]

Since there may be many distinct axiomatic definitions that all provide the same meanings for sentences of L, it makes no sense to talk of the logical structure of a language *simpliciter*. We can only draw the distinction between logical structure and the analysis of concepts, with respect to a specific implementation of the meaning theory for that language.

This does not imply that there may not be some distinctions that must be made in all meaning theories for a certain natural language, or maybe even for all natural languages. However, the

[13] This suggestion has been made by Davidson himself at various occasions, e.g. Davidson [14, p. 22].

[14] Davidson's hope that such a recursive specification can indeed be given for (significant portions of) a natural language, is fostered by the advances in formal approaches towards linguistics, as for instance due to the work of Chomsky.

[15] See [14, p. 30-1], [15, p. 104], [17, p. 71], and [18, p. 136].

[16] That this is what Davidson had in mind is clear from the quote I gave in Section 2.1.

[17] The connection with communication is further elucidated by noting that a theory of truth is counted as a theory of meaning only if it suffices for interpretation. See Section 2.4.

necessity of such distinctions can only be found out empirically, through the recognition of factual coincidence between interpretative theories.[18]

We nevertheless have some preliminary intuitions as to what aspects of grammar we think are superfluous for successful accounts of interpretation. Davidson gives (2.6) as an example of such an intuition. The example shows that – intuitively – we do not have to specify what is meant by (a use of the word) 'good' in order to be successfully interpreted. As long as object language sentences containing the word 'good' are matched to their correlating meta-language sentences by the use of a finite machinery, nothing is missing as far as communication is concerned.[19]

'Bordot is good' is-T-in-English iff Bordot is good. (2.6)

The idea here is that the inexactness of statements involving aesthetic appreciation or evaluation, as well as the uncertainty regarding attributions of belief, are relevant properties of everyday communication. Therefore, a theory that seeks to model interpretation from the point of view of everyday linguistic use should not engage in distinctions that are required for conceptual analyses, unless conceptual analysis is itself part of the communicative endeavour (as for instance in philosophy).[20]

The distinction thus drawn affects how the T-sentences a truth theory generates look like. Calling to mind the T-schema (2.1) it may be thought that the names that fill in for s must be generated by a recursive theory for natural language syntax, while the expressions that fill in for p should be generated by a recursive theory for natural language semantics. The most promising variant for the latter is often though of as some, sufficiently extended, variant of Montague-grammar.[21] But this implementation may in fact be too stringent for the purpose at hand, i.e. it is likely to be too much tailored to conceptual analysis (this is especially clear from the intensional treatment of *oratio obliqua* sentences), whereas all that is required – from the point of view of elucidating communication – is logical structure.[22]

2.3.2 Meaning holism

In the above we explained what a truth theory adapted for a natural language, in broad outline, looks like. But valid T-sentences only state the *truth conditions* of object language sentences; they do not necessarily give the *meanings* of those object language sentences:

> For purposes of interpretation [...] truth in a T-sentence is not enough. A theory of truth will yield interpretations only if its T-sentences state truth conditions in terms that may

[18] This means that for some aspect of grammar to be counted as an instance of logical structure *simpliciter*, i.e. irrespective of the formalization that is used, we need to quantify over a potential infinity of possible theories.

[19] Another instance of such an intuition regarding the distinction between logical structure and conceptual analysis can be found in Davidson's theory for *oratio obliqua*. See [15, p. 104ff14].

[20] Another way in which the distinction can be felt, is by calling to mind that for communication to succeed, at least with respect to sentences such as (2.6), it is simply not necessary to know (exactly) what another person means by taking something to be good. Again, this fares on the assumption that conceptual elucidation is not itself the topic of the communicative act.

[21] For the original specification of this framework, see [32] and [33].

[22] This is why even the most proficient formulations of truth-theoretic semantics are much less detailed than most Montague-grammars.

be treated as 'giving the meaning' of object language sentences. Our problem is to find constraints on a theory strong enough to guarantee that it can be used for interpretation. [19, p. 150]

The point is that a truth theory that is *globally* licensed for meaning in the above described way, need not *locally* license the inference from (2.7) to (2.8).

$$s \text{ is-}T\text{-in-}L \text{ iff } p. \tag{2.7}$$

$$s \text{ means that } p \text{ in } L. \tag{2.8}$$

Truth conditions, evidently an important aspect of meaning, do not necessarily exhaust meaning. So there will be many different T-sentences for a single object language sentence, all with co-extensive but non-synonymous sentences filling in for p. The veracity of (2.9) illustrates this.

'Snow is white' is-T-in-English iff grass is green. (2.9)

The example shows that we must exclude extensionally equivalent T-sentences, and thus extensionally equivalent truth theories, i.e. extensionally equivalent instantiations of the predicate 'is-T-in-English' that do not give (what we intuitively take to be) the proper meanings of the sentences involved. We shall now illustrate what additional constraints are to be put upon a theory of truth in order to guarantee that it can be used for interpretation, i.e. so that it will not generate T-sentences such as (2.9), and such that it will locally license the inference from (2.7) to (2.8).

Davidson thinks the problem of non-interpretative T-sentences such as (2.9) can be resolved by holding the view that only a *holistic* theory of truth will be interpretative. Even though for individual T-sentences, occurring in isolation, it will not generally be allowed to equate meaning with truth conditions, once we define truth conditions for *all* the sentences in the language *at once*, the truth conditions of the individual T-sentences – now belonging to a complex of other T-sentences – will come to automatically exhaust the meanings of the associated object language sentences.

The purpose of a *holistic theory of truth* is not only to filter out non-interpretative truth theories, but it has an important bearing on interpretative truth theories as well. This can be illustrated with respect to the trivially interpretative T-sentence (2.10). If we suppose that the meta-language encompasses the object language, then (2.10), together with all of its disquotational cousins, would constitute an interpretative truth theory that can be legitimately said to give the meanings of the sentences of the language under consideration.

'Snow is white' is-T-in-English iff snow is white. (2.10)

But from this example it is immediately clear that such disquotational T-sentences are, all of them, otiose with respect to the *elucidation* of the meanings of the sentences involved. The fact that 'Snow is white' means that snow is white does, all by itself, not add anything substantial to our understanding of meaning.[23] Davidson's meaning holism is invoked again to solve this problem. For

[23]So Davidson says that "[t]he theory reveals nothing new about the conditions under which an individual sentence is true". See [14, p. 25].

it is only when we consider the position of individual T-sentences such as (2.10) within the semantic web spun by all the other T-sentences that are generated for the same object language, that our understanding of meaning is enhanced.

The merit of a truth theory that is given for a whole natural language *all at once*, is not that it specifies what the sentences taken in isolation mean, but lies rather in that it tells us how every T-sentence is related – in a particular and potentially unique way – to the totality of T-sentences that follow from the truth theory for that language. Even though snow may be white iff grass is green, it is not very likely that (2.11) or (2.12) will be valid T-sentences in a truth theory that ought to work for the whole of English.

'Snow is cold' is-T-in-English iff grass is cold. (2.11)

'Chalk is white' is-T-in-English iff chalk is white. (2.12)

Since the truth conditions are specified in terms of the satisfaction relation, which pairs n-tuples of objects to subsentential expressions, the extension of the predicates $Snow$ and $White$ must be uniform throughout the truth theory (or else lexical ambiguity must be instituted). It is unclear how such a uniform reading of the satisfaction relation would be possible for (2.10) in combination with (2.11) and/or (2.12).

We now understand how the individual T-sentences ought to be fitted into the semantic web of the total theory. We also understand it to be Davidson's claim that such an interconnected web of recursively specifiable truth conditions constitutes all there is to concept of meaning. The only lacuna left in the present account is an explanation for how it is that we come to know that a certain name in s and a certain meta-linguistic expression in p belong together in a T-sentence.

Once we have a certain amount of names related to their respective meta-linguistic expressions (plus an underlying recursive mechanism that automates the task of drawing these interrelations) it will be fairly easy to match new names to new meta-descriptions. This is possible since the number of co-extensive theories is considerably narrowed down due to the already established specification of the satisfaction predicate. A bigger problem is how the process of theory construction gets going in the first place. Since from the holistic condition that is necessary in order for the theory to be interpretative, it follows that it is only useful to construe a truth theory for a large enough potion of a natural language *all at once*.

In order to elucidate this point, Davidson came up with the thought experiment of the radical interpreter.

2.4 Radical interpretation

Radical interpretation is concerned with the act of interpretation in the absence of a preliminary theory of meaning. It thus focusses upon an interpreter's ability to construe a meaning theory from scratch. The thought experiment of a person who finds herself within a community of speakers of an entirely foreign tongue is meant to give substance to the theoretical requirement that a theory of meaning ought not to resort to the thing it intends to explain, i.e. meaning.

Radical interpretation thus constitutes the reversal of the Tarskian view, according to which we possess the implicit ability to assign meanings to sentences, but do not have a theoretical description

of what it is that we are doing. For now we take the view that we have no initial intuitions regarding meaning to go on, but do have the ability to construe meaning theories that we hope will match factual instances of communication.

The meaning theory that gets construed in the process of radical interpretation must be a theory the knowledge of which would put a person in a position to be able to interpret each of a potential infinity of sentences that belong to the language the meaning theory is taken to be a theory of.[24] [25] [26]

There is always some residue of knowledge that, though not part of the meaning theory, is needed in order to be able to interpret others. This consists of knowledge that can reasonably be attributed to a 'normal' human being, i.e. one that satisfies certain minimal conditions that make it earn the predicate 'rational'. The notion of a meaning theory is always related to the practice of communication, in the sense that success in the latter is to be thought of as the teleological purpose of the former.

This teleological purpose can be further fleshed out as the attainment of a set of axioms that generates all and only T-sentences for the natural language sentences we are defining meaning for. The successful fulfilment of this purpose can only be appreciated by concrete instances of communication. The measure of success for the speaker is that the interpreter is observed to understand his words (for instance in response to commands). The measure of success for the interpreter is whether the T-sentences her theory generates as hypotheses for future instances of interpretation, persevere in the face of empirical evidence.

2.4.1 Meaning, belief, and charity

A complicating factor in the endeavour of radical interpretation is that a sentence a speaker utters is true both because of what the speaker believes and because of what he intends his words to mean. For instance, if a speaker believes that big dogs are dangerous, he will say 'That is a dangerous dog!' (pointing to a big dog), whereas the same statement will not be uttered in the absence of the belief (excluding the possibility of lies and sarcasm for the moment).

In order for an interpreter to understand what the speaker means, it is therefore necessary that she assigns beliefs to the speaker and meanings to his words, both at the same time. How is the interpreter able to make this double assignment, provided that in the sphere of interpretation belief and meaning can not be distinguished?

Interpretation, in order to be possible at all, must allow the interpreter to form an interpretative truth theory on the basis of evidence that is readily available to her. Since a detailed overview of someone's beliefs can be no part of such readily available evidence, it is a theoretical precondition for the possibility of radical interpretation that at the onset of the interpretative endeavour one of the two indistinguishable components – either meaning or belief – must be kept constant so that the other can be studied with respect to instances of communication.

[24] See Davidson [20, p. 172].

[25] All that I am concerned with here is the statement that if a person had explicit knowledge of the theory adumbrated, then she would be able to interpret. Whether the antecedent is true or not is irrelevant for the present purpose.

[26] In practice we will always see a *degree* of interpretability, which runs parallel to the degree to which a certain interpreter can be said to have knowledge of the theory under consideration.

The first suggestion is to keep *meaning* constant. This only makes sense if it can, from a theoretical point of view, be maintained that all sentences an interpreter might utter could mean the same thing. That this presumption flies in the face of our intuitions regarding everyday linguistic exchange ought not to bother us in the first place. But what is detrimental to this supposition is that it provides no possibility for the satisfaction predicate to bear out systematic interrelations that only come about once the same atomic parts get shared across *multiple* T-sentences. If all object language sentences are tied to a single meta-linguistic description, it is clear what paucity of holistic structure is derived.

The alternative suggestion is to keep *belief* constant and make meaning vary. That this presumption flies in the face of the fact that we cannot possibly know what beliefs we must attribute to a speaker in advance of understanding what he says, ought not to bother us in the first place. What does induce hope for a way out under this presumption, is that we can – to a certain extent at least – predict what any rational agent would believe relative to a concrete situation. For the shared situation, unlike the speaker's beliefs, is to be counted as evidence that is readily available to the interpreter.

But what is even more important is that the speaker, under the assumption that he answers some marginal notion of rationality, knows all of these things. The speaker knows that the interpreter does not have unmediated access to his beliefs. He also knows that she does observe the same surroundings, and does recognize the same events taking place as he does.

Therefore, in order for a truth theory to be empirically verifiable by a radical interpreter, we need to invoke the *principle of charity*. This principle states that human beings (generally) entertain the same beliefs when confronted with the same situations, which neatly covers the above uncovered assumption. From this it follows that the words the speaker utters at a certain point in time get correlated with aspects of the situation, conditional on the beliefs the interpreter has (or, as interpretation progresses, conditional on the beliefs the interpreter thinks the speaker has[27]).

2.4.2 Speaker veracity

But the principle of charity, as described in Section 2.4.1, is not yet enough. Successful interpretation now seems to depend unduly upon a concession on the side of the speaker to speak in accordance with the shared situation, i.e. (as far as aspects of the shared situation are involved) the truth. What we must add, then, is what I will call the *principle of speaker veracity*.[28]

The occurrence of sarcasm, but also the telling of lies and the making of blunt mistakes, on the part of the speaker, all threaten the interpretative endeavour, even if the principle of charity is invoked.

Let us take a situation in which a speaker is telling a lie to our radical interpreter (similar considerations will apply to similar quirks in the interpretative endeavour, such as sarcasm and making

[27] As interpretation progresses the interpreter uncovers, to an ever increasing extent, both what the words of the speaker (as used by that speaker) mean, and what that speaker beliefs. Since in later stages of interpretation the interpreter is able to (partially) disseminate what the speaker means from what he beliefs, she can attribute further beliefs to the speaker, reasoning from what she already takes him to believe.

[28] Since I treat of speaker veracity as a theoretical assumption, I do not need to look for a 'hold true-attitude' that can be recognized based on more basic evidence, such as behavioural data, as e.g. in Lepore and Ludwig [28, p. 157-8].

mistakes). Provided the interpreter knows, as of yet, relatively little of the speaker's language, what is it that an untruthful utterance made by the speaker consists of?[29]

The conceptual content of a lie is not (only) that the speaker says a thing that he thinks is factually false, but also encompasses that an interpreter is ready to accept as a truth the thing the speaker thinks is actually false, and knows himself to be representing as true with the idea of letting the interpreter believe this.

But this implies that in order for the interpreter to lie successfully, the interpreter must already know the meaning of the speaker's words (or at least the speaker must have reason to believe that the interpreter is in a position to do so). If the interpreter, however, is not in a position to recognize the contents of what the speaker says, then the conditions for the communicative success of the act of lying are not met. In such cases the speaker's utterance is not an instance of communication at all.

The principle of speaker veracity is thus not only a necessary assumption in order to get interpretation going in the first place, but in the early stages of interpretation the only thing that can be conveyed at all is, by and large, truths. In the process of radical interpretation the interpreter is therefore entitled to recognize the speaker as uttering something he holds to be true, conditional on the assumption that the speaker is communicating something in the first place. This is not due to some form of external evidence the interpreter might gain in the light of additional data (e.g. something extrapolated from the speaker's behaviour); rather, it is due to the fact that early on in the process of interpretation the conceptual requirements for the making of lies, mistakes, and the employment of sarcasm are all absent, so that *the stance of holding true cannot but be assumed to apply universally.*

The interlocking chain of belief and meaning is broken into by assuming whatever a speaker says to be true. This means that unless we have evidence to think otherwise, the attitude of holding true is assumed to be the intentional state of the speaker. This assumption, together with the principle of charity, means that unless a radical interpreter has evidence that indicates matters to be otherwise, what the speaker says will be just the thing the interpreter believes when confronted with the same situation.

Early on in the process of interpretation, saying one does not wear red socks when in fact one is, will not effectively provide the interpreter with the mistaken believe that the speaker's socks are green or blue. Instead, it will hamper the interpreter in coming to understand sentences containing the words 'red', 'sock', and 'wear'.

2.5 Indexical elements and demonstratives

Now there is a big problem with the above described account of using a truth theory adapted for a natural language to do the job of a meaning theory. It is the problem this essay is concerned with, namely that of the contextual dependency of the meanings of many (if not all) natural language sentences. For as soon as one attempts to construe T-sentences for such object language sentences as 'You are not to be trusted' and 'Tomorrow it will be too late', it is clear that in order to give

[29]Since we are concerned with communication, we shall focus upon the communicative aspect of lying (and on the communicative aspect of the other quirks of language too).

their truth conditions, we must depend upon non-linguistic information (e.g. who the audience is, and what is the day at which the utterance is being made).

In traditional analyses, typically more oriented towards formal than towards natural languages, such indexical items are dismissed as a non-necessary aspect of natural language sloppiness. The idea is that every indexical occurrence and each demonstrative use can in principle be replaced by a context independent item (e.g. a constant), and that the topic therefore does not necessitate a substantial change to the traditional picture of meaning. But according to Davidson, indexicality and demonstrativity are not that easily discarded as far as a theory of truth for natural languages is concerned:[30]

> No logical errors result if we simply treat demonstratives as constants; neither do any problems arise for giving a semantic truth definition. "'I am wise' is true if and only if I am wise", with its bland ignoring of the demonstrative element in 'I' comes off the assembly line along with "'Socrates is wise' is true if and only if Socrates is wise" with its bland indifference to the demonstrative element in 'is wise' (the tense).
>
> What suffers in this treatment of demonstratives is not the definition of a truth predicate, but the plausibility of the claim that what has been defined is truth. For this claim is acceptable only if the speaker and circumstances of utterance of each sentence mentioned in the definition is matched by the speaker and circumstances of utterance of the truth definition itself. It could also be fairly pointed out that part of understanding demonstratives is knowing the rules by which they adjust their reference to circumstance; assimilating demonstratives to constant terms obliterates this feature. [14, p. 33-4]

It is true that it is not difficult to define the truth predicate so as to treat indexicals and demonstratives as constants. As long as there is a finite number of words in the language that have a indexical and/or demonstrative role, and this indeed seems to be the case, there is no complicating factor as far as truth theory construction is concerned. Then the claim is that there would be a problem regarding 'the plausibility of the claim that what has been defined is truth'.

But the rest of the quote is misguided. The problem regarding the plausibility of the claim that what has been defined is truth, is actually the problem that the T-sentences that are thus construed are no longer in line with our primordial intuitions regarding the concept of truth. So what is claimed to be at stake here by Davidson, is actually Tarski's adequacy criterion (see Section 2.2).

The 'Aristotelian' conception of truth, that was fundamental to Tarski's theory (while assuming the translation between object and meta-language to be a trifle), has in the theory that is adapted for natural languages been replaced by what I have called the principle of speaker veracity (see Section 2.4.2). This means that once we adapt the theory for natural languages, truth is assumed to be a trifle, and the theory's adequacy is now defined in terms of translation/interpretation.

[30] It is unclear to what an extent Davidson thought of indexicality as a necessary element of a natural language. But it is telling that when he attempts to explain his holistic conception of meaning, he is invariably making use of sentences that contain indexical elements:

> [(2.9)] could not belong to any reasonably simple theory that also gave the right truth conditions for 'That is snow' and 'This is white'. [14, p. 26ff10] (footnote added in 1982).

Fodor and Lepore [23, Ch. 3] claim that indexical elements are indeed necessary for Davidson's meaning holism.

All these things considered, Davidson has the matter reversed (and wrongly so) when he says that what suffers (under the demonstratives-as-constants approach) is "the plausibility that what has been defined is truth." (Actually, this is what Tarski would have said.) Instead, what suffers in our case is *the plausibility that what has been defined allows a rational agent to interpret another.*

The right question is therefore *whether radical interpretation requires a special treatment for indexical elements*?[31]

2.5.1 From sentences to utterances

We can alter our truth theory so that it works for utterances (rather than sentences taken exempt from use). The problem here is to determine what characteristics of the context of utterance are relevant with respect to the process of theory-refinement that makes a truth theory interpretative. A quick skim over instances of natural language indexicality suggests that at least an appeal to the

[31] As an instance of how the notion of communicative success in radical interpretation can arbitrate over the problem of indexicality, we take the word 'you'. Concrete instances of communication will invalidate an account in which the word gets interpreted with the meta-linguistic 'you'. The reason for this is that the extension of the word 'you' in the context in which the object language sentences is uttered, potentially differs from its extension in the context in which the meta language sentence is used.

We can explain in some more detail how, with respect to the word 'you', the relativization to speaker (and time) would be introduced during the act of interpretation. Suppose the speaker says 'You are wearing a red shirt' while looking at A, and A is wearing a red shirt. In this case a constant that refers to A suffices. The speaker says the same thing about B, true again. Now 'you' might mean existential quantification. Now the speaker says 'You are not wearing a red shirt' to C, who is wearing a green shirt, true again. Now the existential quantification interpretation is no longer the strongest interpretation; this is now (2.13).

'You are wearing a red shirt' is T iff the person the speaker is talking to is wearing a red shirt. (2.13)

Now the phrase 'the speaker' – as it occurs in (2.13)'s RHS (right hand side) – is not exact enough, since it leaves out the fact that the described speaker is the speaker of the object language sentence that features in the LHS (left hand side).

'You are wearing a red shirt' is T iff the person the speaker of 'You are wearing a red shirt' is talking to is wearing a red shirt. (2.14)

The reference to the object language sentence from within the T-sentence does not lead to a contradiction (as in paradoxes), but it fails to designate the right person. For there may be many speakers of the same sentence. This can be solved, since for each instance of use the speaker can be uniquely lifted out while still being related to the same object language sentence, as in (2.15).

'You are wearing a red shirt' is true in English iff the person Sherlock Holmes is talking to is wearing a red shirt while Sherlock Holmes is saying 'You are wearing a red shirt'. (2.15)

Convention has it that the interrelation of speaker and object language sentence is located within the LHS (because then the name for the object language sentence only has to occur once). We can generalize this for all speakers, as in (2.16).

(S) 'You are wearing a red shirt' when spoken by S is T iff the person S is speaking to is wearing a red shirt. (2.16)

The way in which I shall write this, including tense, is (2.17).

'You are wearing a red shirt' is $T(S, t)$ iff the person S is speaking to is wearing a red shirt at t. (2.17)

speaker of an utterance and to the time at which the utterance was made, ought to be incorporated.[32] [33]

In order to incorporate context-sensitivity, we alter the T-schema from (2.7) into (2.19).[34] [35]

$$(S)(t)\ s \text{ as spoken by } S \text{ at } t \text{ is-T-in-}L \text{ iff } p(S,t) \tag{2.19}$$

Relativizing the truth predicate to speaker and time contributes to the construction of a correct truth theory. For instance, we are now in a position to rule out T-sentences such as (2.9), because they are to be generated by making use of the same satisfaction predicate that is used for T-sentences in which indexicality occurs in the LHS (left hand side), as e.g. (2.20) shows.[36]

$$\text{'This is white' is } T(S,t) \text{ iff what is demonstrated by } S \text{ at } t \text{ is white at } t. \tag{2.20}$$

Relativizing the truth predicate to contexts of use does not mean that we now have defined a truth theory for a specific language user. That a sentence s, containing a personal pronoun, has a different meaning when uttered by different speakers, is a property of a single language (e.g. a property of English in the case of (2.20)).

2.5.2 Context dependence and radical interpretation

In radical interpretation the contextual dependency of sentences is not a threat that tries to crack the shell of shiny context independent meaning. Rather, context dependent terms provide an opportunity for the radical interpreter's relating the semantic theory she is trying to develop to empirical data that is plausibly available to her:

> Sentences with demonstratives obviously yield a very sensitive test of the correctness of a theory of meaning, and constitute the most direct link between language and the recurrent macroscopic objects of human interest and attention. [14, p. 34-5]

[32] See Davidson [14, p. 33-5].

[33] Another, equally important parameter seems to be the hearer, i.e. the one to whom the utterance is intentionally directed. (Davidson makes this suggestion in [16, p. 58].) But since we will assume the standpoint of radical interpretation, this parameter will often be instantiated by the interpreter; it will thus remain constant throughout (most of) the interpretative effort. (It might still be introduced for observed instances in which the speaker directs his utterance to some third person.) It is not clear whether this parameter – implemented in this restricted sense – will be helpful in narrowing down truthful theories to interpretive ones.

[34] The speaker S and the time t can be used in specifying the truth conditions in p, as long as they do not enter an opaque context.

[35] Lepore and Ludwig [29, Ch. 4] show that this formulation is not quite correct. The words 'understood as if spoken by' must replace the words 'spoken by'. The alteration is needed in order to rule out problematic cases like (2.18). This T-sentence would be incorrect for someone who has been the king of France, but who has himself never made that utterance.

$$(S)(t)\ \text{'Je suis le roi de France' is-T-in-French}(S,t) \text{ iff } S \text{ is the king of France at } t. \tag{2.18}$$

[36] If the language in which the object expression is stated is evident, i.e. most of the time this is English, I will leave the name of that language out of the name for the truth predicate.

Davidson takes the identification of context-sensitive items to constitute the second phase in a radically interpretative undertaking (the first phase is to uncover the logical constants and quantificational structure, in order to get the tautological and contradictory utterances in the language right), thus uncovering the truth-conditions for those contingent sentences whose contingency can be directly related to aspects of context.

Why is the truth theory relativized with respect to speaker and time? Because the claim is that with respect to every instance of communication, any interpreter will always come up with a truth predicate that is relativized to both speaker and time (and nothing else) as the optimal interpretative device to handle that specific instance of communication. Observe that, due to the way in which the problem is now formulated, there can be no *linguistic* argument that makes us decide as to what is a contextual parameter and what is not. In the determination of parameters the only thing that is important is that the hearer can define rules that fit into the compositional structure of her axiomatization (or that the semanticist is in a position to do this for her), so that the T-sentences the theory generates are in closest agreement with the utterances of the speaker on factual instances of communication.

Now maybe in such concrete instances of communication it always happens to be the case that relativizing to speaker and time suffices in order to make interpretation succeed. But while a theoretical possibility, this would be an entirely accidental outcome. A first indication that speaker and time might not suffice in order to come up with a successful end result, is given by Austin's well-known example (2.21).[37]

France is hexagonal. (2.21)

The truth conditions of (2.21) differ relative to the use to which the sentence is being put. If we use the term 'hexagonal' in a book on mathematical shapes, then rather stringent conditions regarding the shape of France apply in order to make the sentence come out true. If we use the term in a practice in which we are to give rough characterizations of countries, so that people might easily identify them on a map, different (i.e. looser) conditions regarding the shape of France apply in order for (2.21) to come out as true.

Another indication that the contextual aspects to which a truth predicate ought to be relativized need not be universally the same, is given by observations regarding the vast differences that exist between speakers, hearers, and languages. For instance Tzeltal speakers of Tenejapa (a province in highland Chiapas, Mexico) use words that refer to aspects of the abstracted landscape plane in order to give directions, e.g. 'uphill' and 'downhill'. Other speaker, me for instance, use words relative to an egocentric coordinate system, e.g. 'left' and 'right'. There are other speakers of English who use an absolute frame of reference instead, e.g. 'north' and 'south'. [38] It is clear that in these cases different aspects of context ought to be attended to in order for an interpreter to come out with the right T-sentences.

We conclude this chapter by nothing that it is easy to see that both the speaker and the time of utterance are important parameters for a vast number of sample sentences involving indexicality

[37] It occurs in Austin [1, p. 143].

[38] These data are taken from Levinson [31].

and demonstratives. But whether these two parameters suffice, and if they do why they suffice for narrowing the set of extensionally equivalent truth theories down to only interpretive ones, is not clear at all. The unclarity regarding the sufficiency/insufficiency of speaker and time as the only two contextual parameters against which to relativize a truth predicate, constitutes the core of the minimalism/contextualism-debate. The topic has not been couched in terms of radical interpretation though. Therefore, we will first give an exposition of the discussion in the next chapter. Only thereafter, in Chapter 4, will we be in a position to formulate the solution to the problem of the context sensitivity of meaning in terms of truth-theoretic semantics.

CHAPTER 3

THE CONTEXTUALISM/MINIMALISM DEBATE

3.1 Overview

In the present chapter I give an overview of the contextualism/minimalism-debate to date. In Section 3.2, I give an overview of the conceptual distinctions that have been central to the debate. I begin by discussing quantitative delimitations of context sensitivity in Section 3.2.1; then the distinction between overt and hidden items will follow, in Section 3.2.2; after that the distinction between subjective and objective aspects of context is treated of in Section 3.2.3; the distinction between top-down and bottom-up processed of contextual supplementation will be the topic of Section 3.2.4; and finally, the obscure stance that I term 'perspectival contextualism' will be discussed in Section 3.2.5.

After the conceptual distinctions have been discussed, in Section 3.3 I turn to intuitions that have been claimed to be aroused by specific examples. I will discuss the ways in which these intuitions have been claimed to have consequences for a theory of meaning. First, in Section 3.3.1, I will treat of intuitions that arise when the same expression in used in different situations. Thereafter, in Section 3.3.2, I discuss intuitions as to semantic underdeterminacy, i.e. the feeling that the semantic content that traditional theories provide is somehow not 'propositional' in nature. In line with the discussion of these incompleteness intuitions, I picture the broader problem of background assumptions in Section 3.3.3. The last section that is concerned with intuitions regarding context sensitivity, i.e. Section 3.3.4, will deal with the possibility of moderate contextualism.

After these intuitions have treated, I will end this chapter with Section 3.4, in which I discuss the compositionality criterion, accommodated so as to handle instances of contextual dependency. It is in terms of that notion that I will then define the positions of minimalism, radical contextualism, and my own position: truth-theoretic contextualism (that is described in Chapter 4), in Section 3.4.1.

3.2 Radical contextualism versus minimalism: Conceptual considerations

3.2.1 Quantitative delimitations of context-sensitivity

The first conceptual aspect that plays an important role in the minimalism/contextualism-debate revolves around the extent in which linguistic meaning ought to be established relative to context. The discussion has, in this respect, been oriented towards the more narrow question of *how many expressions* of a certain language, e.g. English, require an appeal to context in order to establish their contribution to the meanings of the sentences in which these expressions occur.

The authority in this respect is Kaplan [25]. According to him only a very limited number of expressions is context-sensitive: "[...] the personal pronouns 'I', 'my', 'you', 'he', 'she', 'it', the demonstrative pronouns 'that' and 'this', the adverbs 'here', 'now', 'tomorrow' and 'yesterday', and the adjectives 'actual', 'present', and others."[1]

The most important representatives of semantic minimalism, Cappelen and Lepore, take over Kaplan's list of context dependent expressions, but make some minor additions. They add the so-called contextuals, which consist of two groups: the contextual nouns like 'enemy', 'outsider', 'foreigner', 'alien', 'immigrant', 'friend', and 'native', and the common adjectives like 'foreign', 'local', 'domestic', 'national', 'imported', and 'exported'. They also add the words 'ago' and 'hence(forth)' as context sensitive adverbs.[2] They call this set the *basic set*.[3]

Even though this has been the customary way of making the delimitation, there are different ways of delimiting the extent in which meaning is to depend on context, i.e. ones that do not involve the identification of a limited set of natural language expressions. In Section 2.5.2 we saw that Davidson imposed a quantitative limit on the number of *contextual parameters* a theory of meaning ought to invoke (according to him, these are speaker and time of utterance, see Section 2.5.1).[4]

There seems to be no argument as to why the basic set ought to provide the right way of quantitatively delimiting the context dependency of meaning. In Section 2.5.2 I already argued that neither is there such an argument for Davidson's quantitative restriction on the contextual parameter set. I therefore conclude that the quantitative delimitation of the context sensitivity of meaning has, as of yet, not been sufficiently argued for.[5]

Regardless of the unfoundedness of the quantitative delimitation, it is often taken to be central to the demarcation between the viewpoints of minimalism and contextualism[6]:

> The most salient feature of Semantic Minimalism is that it recognizes few context sensitive expressions [i.e. those in the basic set, plus or minus a bit], and, hence, acknowledges a very limited effect of the context of utterance on the semantic content of an utterance. [10, p. 2]

There is something wrong with this passage. From the fact that only few expressions are context sensitive it does not follow that there is a limited effect of context on semantic content. For it surely is a theoretical possibility that there are only very few context sensitive expressions (e.g. a single one), but that their semantic content depends upon a vast number of contextual features (e.g. the entire world-state). Whereas, on the other hand, it may be the case that a vast number of expressions

[1] See [25, p. 489].

[2] See [10, p. 1]).

[3] I will be talking about the basic set from now on, indicating either the set originally defined by Kaplan, or some alteration of it; since minor distinctions in the number of words that are included will not affect my argumentation.

[4] I have not been able to find alternative distinctions on the level of contextual aspects/parameters in the recent minimalism/contextualism-debate, though this clearly is a theoretical possibility.

[5] The tests and intuitions that e.g. Cappelen and Lepore provide (and which we will discuss in Section 3.3.1) cannot justify the choice for a specific set of linguistic expressions either. As long as these tests entice conflicting intuitions (and this is what they do) they can not provide independent reasons as to why a specific set of expressions ought to stand out.

[6] See also Borg [7, p. 350], where the same distinction is made in a similar way

is context sensitive (e.g. all the expressions in the language), yet the number of contextual features that impinge on the meanings of all those expressions is, at the same time, quite limited. This would make certain strands of radical contextualism less context sensitive than certain strands of minimalism![7]

Unless an additional argument is supplied, the restriction of context-sensitivity to expressions from the basic set, as it stands, must be removed from the definition of minimalism.

3.2.2 Overt/hidden context-sensitive items

Oftentimes, a distinction is drawn between context-sensitive items that are *overt* and those that are *hidden*. Unsurprisingly, provided the unwieldy nature of natural languages, the way in which the boundary is to be drawn in not a trivial matter. For every distinction that one might come up with there will be ample examples that do not fit in with either of the thus distinguished sides.

First off, if we propose to define overt items to be those visually presented within the surface linguistic string, we ought to count out structures elided because of syntactic ellipsis, like e.g. the VP in (3.1).

$$\begin{aligned} &A\text{: Has John wiped the carpet and bought the roses?} \\ &B\text{: Yes he has [wiped the carpet and bought the roses].} \end{aligned} \tag{3.1}$$

A slightly more fruitful definition would be to call an item overt iff it is present within the 'immediate linguistic environment of the utterance'.[8] But this makes the inaccuracy of the distinction depend upon a preliminary specification of what it is for a linguistic item to belong to the immediate linguistic environment of the utterance, i.e. the original problem is traded in for a new one.

Yet another attempt is to call an item overt iff it appears within the syntactic deep structure of the uttered sentence's parse tree (thus giving more substance to the elusive phrase 'immediate linguistic environment'). A problem with this – somewhat better – definition is that the distinction between overt and hidden items is now taken to be conditional upon the syntactic analysis that is employed.

If this third definition still does not seem satisfactory, one ought to keep in mind that the choice over syntactic structure is to be taken independent of any of the semantic properties of sentences. This means that although the notions of overtness and hiddenness will vary, this variation will always have a syntactic origin and will thus not be unduly blurred by semantic considerations. However the case may be, it is clear that we cannot say of an arbitrarily item whether it is an overt constituent or not, since we can only do so conditional on the syntactic theory that is employed.

[7] I want to point out that my critique of using the demarcation of the basic set in order to define minimalism is not identical to Borg's critique of Cappelen and Lepore's division of radical and moderate contextualism. In [7] she argues that a contextualist can hold that there are forms of context sensitivity that are not captured on the model of the context sensitivity of members of the basic set (i.e. the indexical model).

Though I think Borg would have a good point here if she would have come up with samples of different, i.e. non-indexical, models of context dependence, what I take to be problematic about Cappelen and Lepore's treatment of the basic set is their flawed argumentation. It leads from a quantitative assumption to a qualitative conclusion. (And this is the same mistake we find in Borg [7] too.)

[8] The formulation occurs in Borg [6, p. 213].

Yet the condition of the syntactic theory that is being used is never incorporated in discussions regarding the distinction between overt and hidden items. From this it follows that differences in opinion regarding this topic may as well be due to implicit differences regarding the syntactic structure of the sentences involved.

(3.2) is an example that has been often used in the literature in order to illustrate the notion of a hidden item (e.g. [36, p. 206], [42, p. 9], [43, p. 176], [49, p. 415]).

It's raining. (3.2)

Contextualists maintain that it is evident (in some sense of the word) that the location of the raining event is part of the meaning of (3.2). Let us suppose (3.2) is uttered in Amsterdam, then the meaning of the sentence would not simply be *that it is raining*, but *that it is raining in Amsterdam*. Some contextualists maintain that this is so because the verb 'rain' has an unarticulated locational argument.[9] Minimalists do not allow the verb to be analysed in this way (i.e. with an argument position for the location of the event); they claim that the sentence simply means *that it is raining*, regardless of any further intuitions regarding that meaning's incompleteness.

Now many an argument supporting either of both sides can be introduced at this spot, but we will circumvent a rather tedious discussion by noting that whether or not the verb 'rain' contains an argument for the location of the rain-event, this does not prove the term to be context dependent. It may still be possible, for instance, that the additional argument is a variable that is (existentially) being quantified over.[10] So the example actually makes a very weak case for contextualism. Let us add, then, the assurance that the additional argument in the logical form of the verb is to be instantiated relative to the context in which the sentence is uttered. Now we have contextual dependency proper.

As a matter of fact, semanticists are often quite undisturbed by hidden syntactic items in matters that have nothing to do with context sensitivity. We call to mind instances of syntactic ellipsis, postulations of trace variables in the analysis of questions, and analyses of sentences such as 'John has [at least] three children', according to which the qualification 'at most three children' must be added.

Moreover, the trouble minimalists seem to have with the hidden argument position in (3.2) has nothing to do with the contextualism/minimalism-debate, since the thing that ought to be contextually filled in at the argument position is the easiest thing to come by in indexical land, i.e. specifying the location relative to the speaker of the utterance, an ingredient the context dependency of which every minimalist concedes to, as the inclusion of the word 'here' in the basic set attests to.

From these considerations I conclude that the overtness/hiddenness-distinction, at least in the way in which it has as of yet been introduced into the debate (i.e. disregarding the underlying syntactic structure of the sentences involved) should not be taken to parallel the minimalism/contextualism-divide.

[9] Not all contextualists agree with this, since the location can also be added due to general contextual considerations, as Recanati [45] maintains. This view coincides with that of minimalists: the locational argument is not part of the semantic content of (3.2), but it is added later on in the pragmatic process.

[10] See StanleySzabo2000a and Recanati [42], but also Borg [6, p. 230] in her treatment of liberal truth-conditions, discussed in Section 3.3.2.

A notable participant of the debate who takes the distinction between overt and hidden terms to be of central importance is Recanati [45]. He maintains that a locational variable is added to (3.2) by a process of contextual supplementation. He claims that the process is triggered by general considerations of a *metaphysical* nature. So the location of the rain-event is specified contextually because every event, metaphysically speaking, is taken to have a location.

But I see no communicative purpose in including metaphysical necessities into meaning. The addition of extensionally quantified locational determinations, though necessary from a metaphysical point of view, often enough seems to convey nothing that is of communicative interest. It seems to be an entirely contingent matter whether some metaphysical necessity is accompanied by its communicatively necessary cousin or not.

Suppose that it is a metaphysical necessity that a specific number of cherumbines can stand on the tip of a needle. Would this influence my everyday use of the word 'needle'? This seems very awkward indeed! Meaning ought not to incorporate metaphysically necessary properties of concepts, even if true, without them being of any use for the person employing the word the concept relates to.

I conclude this section by noting that hidden arguments are not to be added to semantic content because they are metaphysically necessary properties of overt terms. For most situations in which (3.2) is used, it is clear that the locale of the event involved is relevant to the situation of use, and is therefore part of what gets communicated. But there is nothing like a specific conceptual necessity over and above the observation that an item's inclusion is relevant with respect to instances of use.

3.2.3 Subjective/objective aspects of the environment

Another distinction that populates the research literature, concerns the nature of the contextual features that are relevant for retrieving an expression's content in context. On the one hand there are *objective* contextual facts – such as speaker, audience, time, and place of utterance –, while on the other hand there are contextual facts involving the mental state of the speaker (and perhaps relevant mental states of other discourse participants too).

The latter are called *subjective* features of context. To this latter category belong such things as the speaker's intention to speak a certain language (e.g. English), the intended lexical meaning of words in the case of lexical ambiguity and vagueness, what the speaker wants to bring about by his or her speech act (i.e. what is implicated), and what the audience to whom an utterance is directed can be reasonably assumed to know about the discourse topic.[11]

The subjective/objective-distinction often occurs in the definition of semantic minimalism, e.g. in Borg [6, p. 207]. But most of the time minimalism is taken to appeal to both objective and subjective aspects of the environment, as in Cappelen and Lepore [10, p. 148] and Recanati [43, p. 174].

[11]The same distinction is often made under different labels. DeRose [22, p. 916] and Borg [6, p. 29-30] use the terms 'objective' and 'subjective', as I do. Bach [4, p. 72] uses 'narrow' and 'wide' context. Perry [37, p. 595-596] (see also Perry [38, p. 59]) talks about 'automatic' and 'intensional' indexicals. Korta and Perry [27, p. 96] use 'narrow' and 'intentional' to say the same thing. Taylor [53, p. 69, 71] distinguishes Kaplanian contexts that consist of objective features from 'speech situations' that encompass objective contexts, but also include the communicative burdens of the agents involved and the intentional actions they perform. The same distinction, though not with fancy names attached, can already be found in Kaplan [25].

Although there are different opinions as to how the distinction between minimalism and contextualism ought to be defined in terms of the distinction between subjective and objective contextual factors, everyone in the debate seems to agree over the way in which this latter distinction ought to be defined. *Speaker's intentions* are generally seen as the relevant dividing line between objective and subjective features of context. But uniform agreement dissolves as soon as we come to assign specific expressions to either of these categories.[12]

According to Kaplan [25] there are the *pure indexicals*, defined as those items that have conventionally established linguistic rules which are called the expression's *character*. Character alone contains all the information required to secure a referent in a given context of utterance. In the specification of an expression's character, only reference to objective features of the context needs to be made.[13]

The objective features of the context are taken to be finite in number, so that they can be formalized by a finite n-tuple. The meaning of pure indexicals is then determined by rules that assign meanings relative to the objective features that this tuple provides.[14]

The second category of expressions that Kaplan distinguishes, are the *true demonstratives*. They differ from the pure indexicals in that they need a demonstration in addition to an objective specification of their character. A demonstration, however, depends (at least most of the time) on subjective contextual features.[15]

But according to contextualists (and some minimalists) even Kaplan's pure indexicals may not be free of the taint of speaker intentions.[16] Let us take as an instance the words 'here' and 'now'. An appeal to speaker intentions is necessary in order to determine the length of time or stretch of space that is referred to by the speaker.[17] [18]

Thus overviewing the research literature, we see a general trend – shared by both contextualists and minimalists – towards shifting more and more expressions from the objective to the subjective realm. I think the shift has not been carried through far enough, and will argue that *all* of the context-sensitive expressions from the basic set belong to the subjective category.[19]

[12]Instead of categorizing terms as either involving intention or not, it seems more promising to enumerate aspects of communication that require an appeal to such intentions. See e.g. Korta and Perry [27].

[13]See [25] and [26].

[14]To give an example, the meaning of the personal pronoun 'I' is the object that occupies the speaker-position in this n-tuple.

[15]According to Kaplan [25, p. 490ff9], not all demonstrations belong to the subjective category, since there may be so-called *opportune demonstrations*. These are said to 'demonstrate themselves', and are thus independent of speaker intentions. An example of this is the sentence 'Stop that man!' when there is only one man in sight.

[16]See [39], [40], [38, p. 61-62].

[17]See e.g. [43, p. 174].

[18]Kaplan did consider these arguments in [25, p. 491ff12], but took them to be instances of vagueness.

[19]Another way of phrasing this full shift is saying that convention plays no foundational role in determining meaning contextually. See Section 5.4 where the role of conventions in meaning assignment is discussed.

In order to make this clear I shall consider the pronoun 'I', since this expression is the prime example of an objective indexical according to both contextualists and minimalists.[20] If person A says (3.3) to me, and A is indeed wearing a red shirt, then this sentence is easily evaluated to be in line with my meaning theory as of to date. Since the T-sentences of my interpretation theory quantify over speakers, I can instantiate for A and satisfactorily check the validity of the thus drawn hypothesis (3.4).

I am wearing a red shirt. (3.3)

'I am wearing a red shirt' is $T(A, t)$ iff A is wearing a red shirt at t. (3.4)

But unbeknownst to me the person in the red shirt, i.e. A, did not utter (3.3) at all. Instead, it was uttered by a person hidden from sight, call her B (with A accidentally or purposefully miming to the words that B uttered). B might not be wearing a red shirt at all. So I ought to have invalidated my hypothesis, or I ought to have attributed a false belief to B. Instead, I did neither; I saw my hypothesis affirmed and attributed a correct belief to A. It thus turns out to be rather tricky to identify the speaker for an arbitrary context of utterance after all! In the future, I will have to make sure that not only a person's mouth is moving and the speech sounds are coming from the direction of that person, but I will also have to make sure that there is no person hiding behind the one I, perhaps mistakenly, take to be the speaker. Further verification conditions apply when I ponder the possibility that she might be carrying a tape recorder with her, etc.

Or suppose that A is holding a hand doll. An utterance of (3.3) may then be judged true if the shirt the doll wears is red. Does this make the fact whether a speaker is holding a hand doll or not an objective component of the context of utterance? Clearly, the meaning of 'I' is what A intends it to mean. There is no such thing as an objective standard that decides when a person is speaking for himself, or is involved in glossolalia, or is ventriloquizing, etc.

Clearly, the number of contextual aspects that I will have to check for is limited only by the intricateness of the intentions of the person who wishes to communicate something to me, tempered by the things that I can legitimately be believed to attend to in uncovering these aspects.[21]

This does not imply that we ought to grow a paranoia attitude in looking for arcane situations in everyday speech so as to potentially uncover a hidden speaker. I just want to illustrate that the *objectivity* of speaker identity, as opposed to the *subjectivity* of the identity of e.g. the thing demonstrated, is a matter of degree, and not one of categorical distinction. What the word 'I' refers to depends as much on the speaker's intentions as any other word does. Even though a speaker who

[20] Even Recanati, who is an important proponent of narrowing down the set of expressions that exclusively depend on objective aspects of the context, holds this view: "To be sure, the reference of a pure indexical like 'I' is determined by a linguistic rule: the rule that 'I' refers to the speaker." [43, p. 173]

[21] The latter condition is important too, since most of the instances we can come up with in this regard are rather unrealistic and require too much effort from the interpreter. We therefore have to keep in mind that a speaker can only communicate what an interpreter may be likely to catch up on, i.e. our communication principle.

is using the word 'I' often intends to refer to herself, this is not, conceptually speaking, a necessity.[22] [23] [24]

I conclude this section by noting that, according to the way in which these terms have been defined, all contextual features can – at least in theory – be called 'subjective'. It is always possible to talk about 'objective features' under the assumption that certain conditions apply and that the speaker is using words in a 'normal' or 'standard' way. But 'conformity to the norm' is not the same as 'objective'. What is detrimental to the distinction between subjective and objective features of context, even if we define the latter notion in terms of standard or conventional use, is that we can form conventions regarding *every* expression in a language (i.e. this does not lift out a specific set).

The here presented critique does not invalidate the kind of formal analyses that we find in e.g. Kaplan, but it does require us to make some distinctions to the methodology that is involved. We must drop the theoretical distinction between subjective and objective terms (and the distinction between rigid designators and direct referring terms, and the notion of the contingent apriori too). There is no such thing as an n-tuple that represents the objective features for each and every context. In the light of this, the minimalist dream of a finite set of variables against which context dependence can be calculated evaporates.

But this does not mean that the radical contextualist is right in doing away with n-tuples altogether. The only problem is that there is not *one* n-tuple that works for *all* of language. The analyses that involve n-tuples and that are given by Kaplan, once the metaphysical absurdities are brushed aside, are mostly right. Only, *the required n-tuple is conditional on the linguistic task at hand.*[25]

3.2.4 Top-down/bottom-up saturation

The fourth conceptual distinction that we want to discuss is a distinction in the way in which contextual aspects are supplemented, namely whether or not this process is exclusively guided by syntax.

[22] Unless convention dictates that only those utterances in which the speaker intends to convey herself as the denotation of 'I', of course. But acts of ventriloquism are legitimate linguistic uses. Irrespective of such considerations, I will criticize the use of convention in establishing linguistic meaning in Section 5.4.

[23] From this it follows that there is no ground for the distinction between rigid designators and direct referring terms, that we find in Kaplan [25]. Nor is there any ground for his identification of contingent apriori statements.

Contingent apriori statements are defined by Kaplan [25] as follows: If an indexical D specifies that its referent in context U is whatever object satisfies conditions C in U, then we can generate the apriori truth 'D satisfies conditions C.'

So the truth of an utterance of 'I am the person who utters this sentence' is known in an apriori way. (It is not necessary since there are possible worlds in which I do not exist.)

Cappelen and Lepore [10, p. 78-9] use the Kaplanian notion of the non-necessary apriori in order to show that such statements cannot be made for sentences that do not contain expressions from the basic set. Examples include 'Everyone is in the contextually salient domain' and 'Some ducks are only in nonsalient domains'. It is clear that if I am right in my treatment of the subjective/objective-distinction, these arguments must be invalid.

[24] Here I have not included those authors whose opinion regarding this matter I simply do not understand. E.g I do not understand how the linguistic meaning of 'we' "forces the agent to be part of its semantic value", as Corazza and Dokic [13, p. 171ff3] maintains. What kind of 'force' is at work here (for I don't feel it)? Another example: the linguistic meaning of 'he' and 'she' "suggests that the value is also determined by the accompanying demonstration or directing intention" (also from [13, p. 171ff3]). How does a word do such a thing?

[25] How this ought to be implemented will be explained in Section 4.7.

There is uniform agreement, between both minimalists and contextualists, that the latter's position is to be defined as making use of pragmatic processes in a *top-down* fashion. Contextual supplementation is said to be top-down whenever aspects of context are appealed to by *general considerations*. This means that it is not be possible to describe these aspects in a finite way, i.e. by making use of a compositional meaning theory (which is assumed to be available for syntax).

Top-down pragmatic processes are contrasted with the minimalist's adherence to *bottom-up* pragmatic processes.[26] These are distinguished from the former, in that they are triggered by a syntactic component that occurs within the original utterance.[27] [28]

The distinction that one thinks must be made between bottom-up and top-down pragmatic processes, because of the appeal to syntax, heavily depends upon the syntactic analysis that is chosen.[29]

Saturation and modulation

Since the distinction between bottom-up and top-down pragmatic influences is never introduced, as it ought to be, relative to a specification of the syntactic analysis that is used, I can only give an impression of some of the leading *examples* that have driven the discussion.[30]

One type of example is called *free enrichment*, in which the context is taken to supplement the semantic meaning of the expression involved. An example of this is (3.5) and (3.6).[31], where (3.5) is thought of as being uttered by a customer, in a café, to the waitress working there.

I would like coffee, please. (3.5)

I would like a cup, filled to about 80%, of liquid coffee served at my table within an acceptable amount of time, please. (3.6)

What items get added by free enrichment depends upon the general aspects of the context, and upon the even more general notion of world knowledge. The notion of *world knowledge* is meant to capture the fact that not everything that is required to be supplemented will belong to the context of utterance. For instance, I must know what it is like to order a cup of coffee in a café, and the waitress to whom I make the utterance must know what it is like to process such an order in the appropriate way, etc. But neither of us is actively pondering all of this information when performing the everyday act of ordering and serving a cup of coffee. Neither can these roles and patterns be derived (solely) from the context of utterance.

[26] See e.g. Borg [6, p. 29].

[27] Bottom-up pragmatic processes are not necessarily triggered by the surface string, since they can also be triggered by items that occur in the syntactic deep structure.

[28] Cappelen and Lepore do not discuss the distinction between bottom-up and top-down pragmatic processes, but their disquotational account of the minimal proposition can be regarded as an instance of the bottom-up approach to minimalism.

[29] The objection is mainly the same as for the distinction between overt and hidden items, see Section 3.2.2.

[30] For an overview of these and other instances of top-down driven pragmatic supplementations of meaning, see Recanati [42].

[31] Taken from Soames [48, p. 78].

Observe that not all instances of world knowledge can be crammed into the lexical meaning of the words 'ordering' and 'coffee', not even with respect to a restricted context of use (e.g. a café). Firstly, there is no end to the full specification of the way in which a customer intends his coffee to be served to him in response to his uttering (3.5), not even on a purported 'standard use'. For instance, our customer will probably want his cup of coffee to be put on his table (and not on someone else's); when it is place on his table, it must not be positioned too far away from him, so that he cannot reach it; but not too near to the ridge of the table either. Also, our customer does not intend the waitress to redress in the kitchen and serve the coffee to him in a dress that has the same colour as the last car that was parked in the parking lot outside, etc.

Secondly, even if the standard use (or some finite number of standard uses) could be completely specified, then it would still be a character trait of language that, provided the right circumstances apply, and provided the right understanding between speaker and hearer exist, an utterance of (3.5) can – in principle – be used to mean *anything*. There is no way in which one can preliminarily specify all the meanings to which a single sentence might be put. Every contextualist agrees with this of course, that is why they introduced the notion of free enrichment in the first place, instead of postulating a richer lexicon. The minimalist also assents to these observations, claiming that most of these aspects of meaning belong to pragmatics instead (keeping the semantics compact).

Another type of top-down (or non-formalizable, or pragmatic) intrusion into determining the truth-conditions of sentences, is called *deferred reference*. An instance of this is (3.7), as uttered by the same waitress, working at the same café. The phrase 'the ham sandwich' is here used dereferentially in order to refer to the person who has ordered a ham sandwich.

The ham sandwich is angry. (3.7)

Even though in deferred reference the new description bears a relation to the original one, e.g. via John's ordering a ham sandwich the proper noun 'John' gets replaced by the definite description 'the ham sandwich', this relation depends upon the contextual situation in such a way that there does not seem to be a general rule for generating such relations in a formal and computationally tractable way. Most uses of deferred reference appeal to general aspects of the context and to world knowledge that speaker and hearer often turn out to share (e.g. that ham sandwiches are not the kind of things that get angry, that people in the café order ham sandwiches all the time, that a person visiting the café – and ordering a ham sandwich – can get angry, etc).[32]

[32]Related to deferred reference is the notion of *predicate transfer*, defined by Nunberg [34]. It is the process that maps the predicate that is standardly attributed to the verb in the input sentence onto another predicate that is related to the former one in the light of the context in which the sentence is uttered. An example of this is (3.9), where the predicate 'parked' that normally takes as argument a car (as e.g. in (3.8)), is replaced by its counterpart that takes a person (who has a car parked somewhere) as argument.

My car is parked out back. (3.8)

I am parked out back. (3.9)

As long as only the (types of) arguments of predicates are replaced, it appears to me that these examples can equally well be described as instances of deferred reference. The same thing can be said about examples in which the predicate itself gets replaced, provided we widen the notion of 'deferred reference' in order to also cover the reference of predicates. A better title would therefore be *deferred extension*. This term is even broad enough to incorporate irony and metaphor as changes in the extensions of sentences.

Recanati's criterion

Now that we have been given some examples, how do we determine – for an arbitrary sentence – whether or not some contextual supplementation is triggered either by the syntax of the sentence or by the context in which the sentence is uttered? In other words, how do we establish whether a pragmatic process is top-down or bottom-up? (Or, how do we distinguish saturation from modulation?)

Recanati [45] claims to have a *test* that tells us whether an expression is linguistically required to be saturated by some aspect taken from context: if there is no (imaginable) example in which the expression is uttered without a specification of a certain contextual item, then the expression involved must be saturated. In order to illustrate this test, Recanati gives an example for the sentence 'It's raining', where the location of the rain event is taken to be the contextual item[33]:

> I can imagine a situation in which rain has become extremely rare and important, and rain detectors have been disposed all over the territory (whatever the territory – possibly the whole Earth). In the imagined scenario, each detector triggers an alarm bell in the Monitoring Room when it detects rain. There is a single bell; the location of the triggering detector is indicated by a light on a board in the Monitoring Room. After weeks of total drought, the bell eventually rings in the Monitoring Room. Hearing it, the weatherman on duty in the adjacent room shouts: 'It's raining!' His utterance is true, iff it is raining (at the time of utterance) in some place or other. [45, p. 127]

Since the example seems to constitute a legitimate utterance of the sentence 'It's raining' even though no specific location is identified, it is concluded that the supplementation of the location of the event is not linguistically required by the word 'rain'. Instead, this supplementation is triggered by general considerations of a top-down nature, or so the argument goes.

However, if intuitions illustrated by made-up examples are to determine the dividing line between items that are saturated in a bottom-up fashion and those that are modulated (e.g. freely enriched) in a top-down fashion, it is clear that there are no items of the former sort whatsoever.

The following example concerns the verb 'arrive'. Recanati [45, p. 128] takes this to be a clear instance of a term that does require locational saturation in a bottom-up sense. But because of the creativity of language it is clear that the test he proposed works for that expression as well:

> I can imagine a situation in which alien John's arrival on earth is eagerly awaited. There are hundreds of airports around the globe, and on any one of them John may choose to land his UFO sometime soon. Since the first contact with an extraterrestrial being will be an occasion of great importance, all airports around the world are connected to a network, sending the message of John's landing to all radio stations across the globe, causing them to play the 'welcome alien John'-song upon receiving that message. After weeks of waiting suddenly the 'welcome alien John'-song is played on the radio. Upon

[33] See also [41, p.317] and [42, p. 9].

hearing it I shout: 'Alien John has arrived!' My utterance is true iff alien John has arrived (at the time of my utterance) in some place or other.[34]

In the absence of a further specification of what the bottom in bottom-up consists of (i.e. what is the syntactic analysis that is being used), and as long as there are no sensible tests that indicate what contextual aspects are linguistically required by a certain expression, I take it that the bottom-up/top-down-distinction is not a theoretically viable discrimination.

3.2.5 Perspectival contextualism

In the preceding sections we have assumed the viewpoint of radical contextualism as embodying the claim that the meanings of the all expressions within a language depend upon aspects from context (and/or world knowledge). Although this is the common opinion in the debate (and the one that most minimalists oppose to), there is an alternative view that (often) goes under the same name. It differs from the former in that it does not claim that an indefinite number of contextual aspects apply to the process of establishing the meaning for every expression in the language, but – instead – that there is no such thing as 'meaning' to begin with.

Borg [6, p. 224], in characterizing this alternative strand of radical contextualism as it occurs in the work of Travis, has termed this view the *perspectival nature of meaning*:

> Travis counsels that we should properly appreciate the *perspectival* nature of both thought and language: there is no such thing as what is meant abstracted from the purposes, or role, to which a thought or utterance is put, there is no such thing as an interpretation which is not an interpretation in a context and for a reason. [6, p. 224]

My problem with the notion of the prespectival nature of meaning is not that it is downright mistaken, but that it mixes mistakes with truths, thus obfuscating the matter. Let us take the above quote by Borg as a starting point. In a sense it is right that there is no such thing as meaning except when placed within a concrete practice in which language is put to use. But this does not imply that it is impossible to come up with formal theories of meaning. For although a meaning theory abstracts away from most of the aspects of the use to which a linguistic expression is being put, the continual verification of this theoretical abstraction (with the purpose of establishing its correctness) makes it an abstraction that is constructed both in tandem with and tied to instances of use.

Another quote from Travis himself clearly expresses what I take to be the mistaken part of what I term *perspectival contextualism*:

> [Perspectival contextualists] deny that (e.g. English) sentences are in the business of being true or false. They are not ([perspectival contextualists] hold) because there is, systematically, no such thing as 'that which a sentence says to be so'. [54, p. 40]

From this quote it is clear that perspectival contextualism embodies an illegitimate inference from the correct antecedent observation that a theory of meaning can only apply to instances of use, to the incorrect consequent ideology that sentences do not conform – within such a use-based

[34] Of course the word 'arrive' might be taken to be ambiguous between a version that takes a locational parameter and one that does not. But this would only label the problem, not solve it. For we would need to establish the conditions under which either of the disambiguated versions of the word applies.

theory – to very specific truth conditions (that can be both formally specified and generated in a computationally tractable manner).

3.3 Radical contextualism versus minimalism: Intuitions

Since we seem to lack reliable conceptual distinctions that we can use to describe how context influences meaning – overt/hidden, subjective/objective, top-down/bottom-up, all turned out to be problematic divisions –, we will now switch to intuitions regarding the context sensitivity of meaning, and see whether these can bring us any closer to a theoretic appreciation of the problematic.

3.3.1 Context shifting intuitions

The first set of intuitions that we will look at, focusses upon intuitions regarding shifts in meaning that occur once a sentence gets uttered in different contexts. The idea is that if intuitions regarding an expression's meaning can be so different across contexts, then its meaning cannot remain constant. I will call intuitions regarding shifts in meaning across contexts *context shifting intuitions*.[35] [36]

Many test have been proposed for establishing whether an expression admits of context sensitive intuitions. An instance of one such a test is the *inter-contextual disquotational indirect report*[37]:

> Take an utterance u of a sentence S by speaker A in context C. An Inter-Contextual Disquotational Indirect Report of u is an utterance u' in a context C' (where $C' \neq C$) of 'A said that S'. [10, p. 88]

An expression may be context-sensitive if some of the sentences in which it occurs fail on this test. The following sample case that purportedly fails on this test is given by Bezuidenhout:

> We're at a county fair picking through a barrel of assorted apples. My son says 'Here's a red one,' and what he says is true if the apple is indeed red. But what counts as being red in this context? For apples, being red generally means having a red skin, which is different from what we normally mean by calling a watermelon, or a leaf, or a star, or hair, red. But even when it is an apple that is in question, other understandings of what it is to call it 'red' are possible, given suitable circumstances. For instance, suppose now that we're sorting through a barrel of apples to find those that have been afflicted with a horrible fungal disease. This fungus grows out from the core and stains the flesh of the apple red. My son slices each apple open and puts the good ones in a cooking pot. The bad one he hands to me. Cutting open an apple he remarks: 'Here's a red one.' What he says is true if the apple has red flesh, even if it also happens to be a Granny Smith apple. [5, p. 107]

[35] Cappelen and Lepore [10] talk about 'context shifting arguments'. But since these consist of nothing more but felt intuitions, I find the word 'argument' not appropriate here.

[36] Another set of intuitions that are much like context shifting intuitions, are what I would call *inappropriateness intuitions*. These involve sentences whose semantic meaning does not alter between contexts, but that somehow seem less appropriate when uttered in certain contexts. See [7, p. 341], [12, p. 26], and [42, p. 23].

[37] For this and similar tests see [10, Ch. 7].

Bezuidenhout [5] uses this example in order to show that the meaning of 'red' changes due to the context in which it occurs. But Cappelen and Lepore [10, p. 92] object to this. Calling the two contexts that feature in Bezuidenhout's example $C1$ and $C2$, they imagine themselves in yet a different context $C3$. They then claim to be in a position to legitimately report on the apple sorting affair by utterances of the sentences (3.10), (3.11), and (3.12).

In $C1$, Anne said that the apple was red. (3.10)

In $C2$, Anne said that the apple was red. (3.11)

In both in C1 and in C2, Anne said that the apple was red. (3.12)

It is easy to agree with these statements, since they indeed seem to be legitimate reports on the event involved. But multiple meaning assignments to the word 'red' are in line with these reports also. Given the data that is available the word 'red' may mean *red on the outside* in (3.10), *red on the inside* in (3.11), and *red in one way or other* in (3.12). That, as Cappelen and Lepore maintain, a uniform meaning assignment to the word 'red' applies here must be argumented for independently.

Even if one would have the same intuitions as Cappelen and Lepore have, then this would still not imply anything about the semantics of the words involved (i.e. the link between intuitions felt and linguistic properties would have to be argumented for). The same is true for Recanati's, or anyone else's, intuitions of course. Even if one were allowed to extrapolate semantic features from felt intuitions, we would still be faced with the problem that broad agreement on the outcomes of tests for context sensitive intuitions are only available for elements that belong to the Kaplan- or basic set. For all the other expressions the evidence such tests provide us with is simply inconclusive.

But there are more intrinsic problems to context shifting intuitions. One of the problems of the tests regarding context sensitivity lies in the specification of inequality conditions between contexts. For context shifting intuitions to arise, it is insufficient to invoke a number of arbitrary contexts. What we need are contexts that differ from one another according to a so-called *relevant parameter*. To give an example, for the personal pronoun 'I' the relevant parameter is the speaker of the utterance involved. Therefore, if two contexts have the same speaker they are not relevantly different with respect to 'I'.[38] The right intuitions only arise provided the contextual parameter that influences the meaning of the expression we are testing context sensitivity for, changes between contexts.

But how does one establish the contextual parameters that supplement (or alter) the meaning of a given expression? These considerations must be further weakened in the light of the objections that were levelled in Section 3.2.3. So, for instance, the same speaker can intend to convey two different contents in two contexts, the one involving ventriloquism, the other not.

Even for the standard indexical expressions it is not clear what the relevant parameters ought to be on each and every occasion of use. It follows that the choice of the relevant parameter itself depends on intuitions one has regarding a certain expression's context sensitivity. The notion of a relevant parameter does there not legitimately feature in the specifications of context sensitivity tests. The contextual parameter that I take to mark a relevant distinction between contexts is likely to depend upon just the same sort of intuitions that make me intuit a word's contextual dependency in the first place.

[38] In other words, the sentence 'I am walking' has a constant meaning across contexts in which the speaker stays the same.

3.3.2 Incompleteness intuitions

Another batch of felt intuitions that has been thought to bear out the distinction between context-sensitive and -insensitive expressions, are those in which the semantic content of S is P, but the world is taken to be neither P nor not-P. These intuitions imply that the semantic content P does not specify truth conditions for S. For instance the sentence 'The rhino is purple' can be claimed to be neither true nor false, since in the world a rhino is neither purple nor non-purple. It is only purple to a certain extent, under specific lightning conditions, etc. Intuitions that adhere to this scheme I will call *incompleteness intuitions*. These intuitions can be described as a felt inability to evaluate the truth-value of an utterance based on solely its sentence meaning.

What supplementations one must make in order to get a complete characterization of an expression's truth conditions, depends on the context of utterance and on knowledge of a broader nature, i.e. so-called background assumptions (these are treated in Section 3.3.3, whereas in this section we will focus on contextual knowledge of a more narrow sense, i.e. knowledge pertaining to the directly available situation of utterance.).

Incompleteness intuitions and context shifting intuitions are not conceptually related. I.e. from the fact that someone has context shifting intuitions we can not (generally) infer that he has incompleteness intuitions (regarding the same expression). E.g. someone could have the feeling that different things get expressed by utterances of the same sentence under relevantly different circumstances, although all of these constitute full propositions. Conversely, if a person takes a sentence's semantic meaning to be non-propositional in nature, then this does not imply that there are alternative ways, relative to context, that allow the formation of propositional content.

An instance of an incompleteness intuition can be given for (3.13). It is claimed that there is no such thing as being strong *simpliciter*; so we need to supplement the semantic content in order for it to be truth-evaluative. In order for the semantic content to reach the level of propositionality, we need to say what it is that steel is strong enough for (e.g. a strong enough shed, a strong enough statue, etc).[39]

Steel is strong enough. (3.13)

Contextualists hold that (3.13) – unless contextually enriched – does not express a complete proposition. They argue that the sentence does not specify what it is with respect to which steel is taken to be strong enough.

If the criterion for a proposition's being (in)complete is that there are possible questions whose answer is not part of the semantic content that is under consideration (like 'Strong enough for what?' in case of (3.13)), then it is clear that no semantic content – however heavily supplemented – will ever reach the level of propositionality. It is clear that we can always come up with yet another question the answer of which is not yet encoded in the proposition we have.[40]

Another suggestion that occurs in the literature is that semantic content is said to be non-propositional (or not truth-evaluable) if it is unclear how one is to *verify* whether the content that is

[39]The example is treated in Bach [2, p. 127], Bach [3, p. 269], Recanati [44, p. 22-3], and Pagin and Pelletier [35, p. 53].

[40]Cappelen and Lepore [10, p. 63] make this objection.

expressed does obtain in a given world or not.[41] Minimalists react by saying that some contextualists mistake truth-conditions for truth-values, who are therefore committed to verificationism and not semantics.[42]

A third problem with these incompleteness intuitions is that we are not given a criterion of what it means for a semantic content to be propositional or not.[43] From a formal point of view we can define the notion of sentencehood as part of the formal specification of the meta-language. Whether incompleteness intuitions arise thus depends on the meta-language we choose to formulate our T-sentences.[44]

Whenever contextualists claim that a sentence has non-propositional semantic content, they must therefore mean that there exists *no* pair consisting of a compositional meaning function and a logical language, with its associated definition of sentencehood, that is able to assign the right semantic contents relative to context. Since the argument ranges over a potential infinity of meaning functions, only general considerations would suffice in order to prove the point. Sadly, no contextualist as of yet attempted to do so.

Needless to say, without a specification of what is and what a proposition is (and what is not), incompleteness intuitions are indicative of nothing whatsoever. The point is not that the notion is itself difficult to define – after all, every ordinary logic defines sentencehood –, it is just that nobody bothers to formulate the matter nicely.

My take on this is that if a notion is thought to be central to a certain topic, and that notion is at the same time not further specified, then either it is a crystal clear intuition – one that is in no need of elaboration at all (a Cartesian insight, say) – or it is the wrong kind of notion with respect to that topic (and we ought to do away with it in our theorizing). I think neither of these answers are correct, and will thus attempt to further specify the notion of propositionality in Section 4.3.

To end this discussion of propositional incompleteness, we will now discuss Borg [6]'s attempt to solve the problem by introducing the notion of *liberal truth-conditions*. According to the theory of liberal truth-conditions, unarticulated constituents are introduced in so far as the inclusion of such elements is 'intuitively compelling' and/or there is 'good empirical evidence' to support such incorporations.[45] An example of such liberal truth-conditions is (3.14).

$$\text{If } u \text{ is an utterance of 'Jane can't continue' in a context } c, \text{ then } u \text{ is } T \text{ iff} \quad \text{Jane can't continue something in } c. \tag{3.14}$$

Borg's analysis hinges on her interpreting (3.14) as positing a context sensitive element in the LHS of the T-sentence. But alternative interpretations are available, namely ones in which the object- and meta-language are different (despite the appearances of the surface strings involved).

[41] See Recanati [43, p. 185].

[42] E.g. see Borg [6, p. 238-9].

[43] Cappelen and Lepore [10, p. 61] have noted the lacuna.

[44] Since sentencehood is a syntactic notion, it is important to note that the semantics – i.e. the recursive assignment of meaning to syntactic structure – must be defined so as to compositionally bestow meaning upon all things syntactically identified as sentences.

[45] See Borg [6, p. 230].

From the point of view of communication, Borg's interpretation of (3.14) is even excluded (and only the alternative interpretations remain). For the only way in which the added qualification 'something' (as it occurs in the RHS) can be thought to play a role in (3.14), is when the person whose meaning theory this T-sentence belongs to finds a purpose for including this qualification in her attempt to interpret the speaker of the object language sentence on the LHS. And this is an indication that the interpreter and interpretee, at least according to the interpreter, are speaking different languages.

Another problem is that according to (3.14) the object language sentence will be true whenever there is something which Jane cannot continue. But there will always be something in the world that Jane might be said to have not continued as of yet. This means that liberal truth-conditions give nigh to necessary truth-conditions for what seem to be contingent object language sentences. Borg notes this problem herself, but thinks this is just part of the minimalist approach towards semantics. The idea seems to be that the minimal proposition cannot express everything, but must depend for its full supplementation on further pragmatic processes; and these will bear out the sentence's contingency. This shifts almost everything that is of importance with respect to the notion of meaning to pragmatics. Both because of the above noted fault and the here noted undesired consequence, it is clear that the notion of liberal truth-conditions – at least for our purposes – is not an option.

3.3.3 Background assumptions

The word 'cut' is unambiguous (says Searle, though dictionaries disagree), yet it makes quite different contributions to the truth-conditions of sentences (3.15) and (3.16).

Bill cut the grass. (3.15)

Sally cut the cake. (3.16)

According to Searle [47] this is so because background assumptions play a role in fixing satisfaction conditions for the VPs of these sentences. This position is really a variant of incompleteness intuitions, since in order to get something genuinely evaluable Searle's case, we need to supplement the meaning of the terms 'cut' in the light of the background assumptions that are involved. The problem with these background assumptions is that there is (claimed to be) an infinite lot of them. This means that supplementations/explanations will never come to an end (unless we preliminarily take an end in view).

From the point of view of truth-theoretic semantics the problem is not that obvious. As long as the word 'cut' means something like 'remove something from something larger by using a sharp implement', then this meta-linguistic description seems to be applicable in T-sentences for both (3.15) and (3.16). That the grass will be cut with a mower (and the cake with a knife) is not necessarily part of the *meaning* that gets expressed by these sentences.

Imagine a person A uttering (3.15) to you. Afterwards you discover that A in fact cut the grass using a pair of scissors. What is strange (or wrong) in this example is not A's initial utterance, but his queer behaviour (under the standard meaning of the word 'cut' as uttered by him).

But things are about to get quite a bit queerer, as Searle considers the application conditions of the imperative (3.17).

Cut the sun. (3.17)

We do not know what it would mean to obey that order. But does this show that the dictionary meaning of 'cut' is insufficient? It only shows that the meaning of 'cut', which must be somewhat like 'remove something from something larger by using a sharp implement', is not applicable to the sun (at least not in a straightforward way), i.e. that it is unclear what in the case of the sun would be the parts to dissect. But I will argue that it is not due to meaning (solely) that the verb 'cut' cannot be used in application to the sun. Besides (but, of course, not instead) meaning we must consider the practice of cutting. Pending an argumentation (which occurs in the next chapter), I will give an example of this: we are able to follow the command (3.17) if 'the sun' refers to the yellow oval cardboard-make sun that functions as a piece of decoration in the situation in which we make the utterance.

In this example the traditional meaning of 'cut' applies unproblematically. The example is characterized by the fact that there are clear application conditions for cutting a cardboard-make sun. Crucially, there must be conditions in which the utterance of a word (in a sentence) makes sense relative to the situation at hand. In such cases either the old meaning theory applies unproblematically, as in the above example, or we will have to change our meaning theory somewhat so that the meaning theory is made to bestow new meanings onto old expressions, according to the newly encountered use.

These rather tentative statements are now illustrated with respect to utterances of (3.18), which is yet another example Searle gave.

The cat lies on the mat. (3.18)

This sentence seems to have truth conditions, i.e. it is true just in case there is a cat and there is a mat and the former is placed on top of the latter. Searle illustrates the 'descriptive element' of (3.18) by drawing a little picture showing a cat positioned over a pictured mat. But now suppose that the cat and the mat are in the depicted relations, only now they are both floating freely in outer space. In such an imaginary situation the scene could be equally well depicted by turning the pictorial representation upside down (since there is no gravitational field relative to which either of the pictured elements can be said to be above the other). Is the cat still on the mat then? And if so, was the earth's gravitational field one of the things depicted in our earth-bound drawing?[46]

It is a mystery why Seale thinks his talk of background assumptions should be introduced in *semantics*. There is simply no reason why the theory of meaning ought to be the place where background assumptions regarding the broader context of our knowledge of the world would have to be taken into account. My take on the matter is that the problem Searle identifies is real, but cannot be solved in semantics. Although the problem cannot be solved there, this is not a problem for the semanticist, since is shifted outside of the semanticist's province by matching object language sentences with meta-language sentences to which the same background assumptions

[46]The appeal to a pictured representation solely serves didactic purposes of course. The same considerations apply to the meta-linguistic expression 'The cat is on the mat', since here there is no reference to a gravitational field either.

apply. Epistemology (which ought to solve the matter of background assumptions) must explain this phenomenon as it occurs in sentences irrespective of whether they belong to either the object language or the meta-language.

A question that still remains is how we know that our meta-linguistic descriptions do indeed capture *all* of the background assumptions that go into the object language expression? Remember that we are only interested in those aspects of the object language sentence that are necessary for (instance of) communication. This means that if there are truth-conditions we do not capture in our meta-language, because of our leaving out some relevant background assumption, then this will either not matter from the point of view of communication, or it will be found out in the act of communication itself. The only thing the RHSs of T-sentences ought to incorporate are those background assumptions that are relevant for interpreting a speaker, but that are not handled by e.g. the meta-linguistic disquotation.

This means that there may be any number of background assumptions that we do not share with the person we are trying to interpret (potentially, though unlikely, there can be an infinite lot of them). But as long as these differences do not come to the fore during our communicative effort, they do not really matter from the point of view of (truth-theoretic) semantics. Now the fact that (3.18) can only be used unproblematically near the surface of the earth is a perfect example of a background assumption that we need not take heed of in semantics. As long as the sentence is not used anywhere but near the surface of the earth, and as long as the sentences of the meta-language comes with the same presuppositions (which is very likely just in case the interpreter is, like us, not in the habit of travelling through space a lot), there is no communicative purpose in making the background assumption explicit. Just in case we do go on space travel tomorrow, and see cats on mats floating about on a regular basis, considerations regarding the optimization of communicative success will decide whether the gravitational field remark will be added to the RHS of the T-sentence for (3.18) or not.[47]

From this extensive treatment of sample sentence (3.18), it is now apparent why there is nothing problematic about there being a potentially infinite number of background assumptions at play in some (or all) natural language sentences. For even if there would be an infinite number of background assumptions at play in an object language sentence, then this may as well be the case for the meta-language sentence translating it. Since a Tarski-style truth theory (adopted in order to apply to natural languages) requires of a meta-language that it be strictly stronger than the object language for which truth is defined (see Section 2.2.1), it is clear that problem regarding background assumptions can ever arise in semantics.

3.3.4 Moderate contextualist intuitions

We have only been discussing radical contextualism as opposed to semantic minimalism. But most philosophers who call themselves contextualists claim they are somewhere in between these outer poles. Their position is under constant attack of both radical contextualists and minimalists. Even though the proponents of either of these 'extreme' views disagree over who is right on matters concerning semantics, they both agree – and sometimes for the very same reasons – that moderate

[47]For the example at hand this means that only if we want to distinguish cats on mats in outer space from their earth-bound counterparts, i.e. only if we find a use in our lives for (occasionally) drawing this distinction, does it make sense to make the proximity to earth explicit.

contextualism is an untenable position. Moderate contextualists, on the other hand, seek to pose a halt to the slippery slope arguments that both minimalists and radical contextualists employ.

Leslie [30] employs context shifting tests in order to show that the set of context sensitive expressions is larger than the basic set, yet does not extend to all the expressions in the language (thus ensuring a possibility for moderate contextualism). is the most intricate and complete account of tests for context sensitive intuitions that I know of, I still have some very fundamental problems with the tests she comes up with.[48] Allow me to quote two of Leslie's sample tests here, since I think the contrast between the two is particularly elucidating.

The first example is meant to show that 'tall' is context-sensitive:

> Seeing as how he measures 6'3", Tom is tall. He plays basketball from time to time, and once he called me from the court, because he was feeling nervous before the game. To reassure him, I said "well, it'll help that you're tall". He replied "are you kidding me? You should look at the guys I'm up against. I'm not tall at all!". He was right (the other guys were approaching 7'!), even though at 6'3", Tom is definitely tall. [30, p. 138]

The second example is meant to show that 'tall for a pregnant giraffe that is standing up straight' is not context-sensitive:

> Georgina is simply not tall for a pregnant giraffe that is standing up straight. No one who saw her would claim that she was. But the other day we were talking about pregnant giraffes that are standing up straight but have just taken a bath. I said "Georgina is tall for a pregnant giraffe that is standing up straight!" What I said was true, because taking a bath shrinks giraffes by a small amount. Of course, now that I am just looking at a dry Georgina, she is definitely not tall for a pregnant giraffe that is standing up straight. [30, p. 139]

Let us call a pregnant giraffe that is standing up straight a 'praffaight', just for ease of expression. In the second example by Leslie there are two contexts at play. We can further specify these contexts as one in which Georgina, herself a praffaight, is compared to praffaights that have not taken a bath and are just a tiny bit taller than Georgina is. We can imagine a second context in which the praffaights to which we are comparing Georgina's tallness have shrunk a bit because they have just taken a bath. But, in this second context, Georgina has herself not taken a bath, so that she is now the tallest of the flock. This second example with the here made up contexts seems to me to be just as context sensitive as Leslie's former example involving Tom's tallness is. And indeed, we would be somewhat startled to hear that the word 'tall' was context sensitive when applied to people (or, more narrowly, baseball players), but not when applied to praffaights!

Since the above two examples are structurally identical, the distinction that Leslie seeks to make must lie in the 'nature' of the entities involved, i.e. it must be a characteristic of a person that his tallness varies with respect to (groups of) other persons; and must be a characteristic of a praffaight that her tallness cannot vary with respect to (groups of) other praffaights.

[48] My critique of the tests that occur in [30] extends to all the other tests for moderate contextualism that I know of.

I do not object that there is a distinction between the first and the second example; the second one surely appears to be a lot weirder.[49] but I claim that this is not a distinction that can be explained by appealing to *linguistic* aspects (at least not solely). It intuitively makes sense to take the attribution of tallness to a person to depend upon the contextually salient set of basketball players, because we are acquainted with a clear use for such a comparison. The praffaight example appears to be so much less intuitive, since we cannot in the world think of a purposeful use for the comparison involved.

Summarizing, I do not see how any portion of intuition regarding context shifts allows us to draw a neat line isolating a limited set of context-sensitive expressions, and yet this is what moderate contextualism requires. What the discussion of the praffaight example shows, is that although there are no clear distinctions between context sensitive and context insensitive *expressions*, there is a distinction between customary and non-customary *uses*.[50]

3.4 The compositionality criterion

In the preceding sections I have described the important conceptual aspects and intuitions that make up the minimalism/contextualism-debate. The outcome of the discussion has been exclusively negative though. In Section 3.2 we saw that none of the conceptual distinction makes sense when put under sufficient scrutiny. And in Section 3.3 we saw that for none of the intuitions regarding contextuality is it apparent what the consequences for (a theory of) meaning ought to be.

But there is one important conceptual distinction that was not discussed in Section 3.2, it is the *compositionality criterion*. The reason why I introduce it here, at the end of this chapter, is that it is the only consideration that can be draw from the debate that does instigate a sensible division between contextualism and minimalism.

The compositionality constraint is one that every formal characterizations of meaning ought to consent to (or so is the claim). This section relies heavily on work by Pagin and Pelletier [35]. They have thought through the alterations that must be made to the traditional formulation of compositionality in order to make it apply to a meaning theory that involves contextual aspects too. They have altered the global formulation of the principle of compositionality in the following way[51]:

Compositionality for context sensitive languages: For any complex expression and context, the meaning-in-context of the complex expression is a function of the meaning-in-context of its parts and the mode of composition.

Furthermore, they have defined two formulae schemata to which composition functions must conform in order to meet the compositionality criterion. The formulation of the first schema makes use of the following concepts: The *general composition function* ρ mapping syntactic operators

[49]The notion of the weirdness of expressions that are not customarily used, e.g. 'Georgina is tall for a pregnant giraffe that is standing up straight' and 'Cut the sun', was discussed in Section 3.3.3.

[50]This observation is crucial for the formulation of my solution to the minimalism/contextualism-debate in Chapter 4.

[51]This is taken from [35, p. 39].

σ onto *particular composition functions* $\rho(\sigma)$.[52] The *meaning function* μ mapping expressions to conceptual structures, that are collated in the set CS. *Conceptual structures* are either unstructured (or atomic) meanings, the set of which is designated by M, or n-tuples that consist of a particular composition function $\rho(\sigma)$ and $n-1$ conceptual sub-structures (also belonging to CS). A context of utterance c is an ordered tuple of formally represented contextual factors. This concludes the elements that belong to the schema for the compositional semantic function.

The semantic function $\rho(\sigma)$, that operates on meanings, must be compositionally in line with the syntactic function σ, that operates on expressions. This is the schema the function that describes the semantics must adhere to in order to be compositional[53]:

Compositional semantic function: $\mu(\sigma(t_1,\ldots,t_n),c) = \langle\rho(\sigma),\mu(t_1,c),\ldots,\mu(t_n,c)\rangle$.

The schema for the compositional pragmatic function adds the following elements: A *modulated evaluation function* μ' which maps conceptual structures onto meanings. The modulation performed is represented by $\mathcal{M}$. This is the schema a function describing the pragmatic process must adhere to in order to be compositional:

Compositional pragmatic function: $\mu'(\mu(t,c),c) = \mathcal{M}(\mu(t,c))$

$$\mu'(\langle\rho(\sigma),\mu(t_1,c),\ldots,\mu(t_n,c)\rangle) = \mathcal{M}(\langle\rho(\sigma),\mu'(\mu(t_1,c),c),\ldots,\mu'(\mu(t_n,c),c)\rangle)$$

According to Pagin and Pelletier, contextual supplementation or modulation can be characterized by functions that adhere to these compositionality schemata. These schemata are neatly defined and all that, but to me it is unclear how a compositional theory that adheres to these principles ought to work *for a concrete example*. Let us take sample sentence (3.19), which requires pragmatic modulation in order to make explicit the intended relation that holds between John and a bat. I assume that (3.19) will turn out to mean something like (3.20), but similar objections can be levelled under alternative interpretations too. We now wonder what is the right modulation $\mathcal{M}$ so that (i) it makes use of aspects we can formalize in a contextual tuple c, and (ii) it makes the second element of the tuple in (3.20) come out as in (3.21).

John's bat is dead. (3.19)

$$\langle\rho(\sigma),\ John,\ \mathcal{M}(\lambda y.\lambda P.\lambda Q.\exists!x,\exists R(R(y,x)\wedge P(x)\wedge Q(x)),\ \lambda x.Bat(x),\ \lambda x.Dead(x))\rangle \quad (3.20)$$

$$\lambda y.\lambda P.\lambda Q.\exists!x.\exists R(R(y,x)\wedge P(x)\wedge Q(x)\wedge R = Own) \quad (3.21)$$

In (3.21) we chose to interpret 'John's bat' as 'the bat that John owns', but alternative modulations of R could have been 'the bat that John loves', 'lives together with', 'got married to', 'is responsible for', 'knows', 'wishes to get influenza from on Friday the thirteenth', etc. Even for this limited set of choices (establishing the full list of potential options here is a task as endless as a person's ability to creatively express herself in language), it is not clear which representation of the

[52]The totality of syntactic operators defines the set of expressions in the following way: a language L is a pair consisting of atomic expressions A and syntactic operators Σ, such that A closed under the operators in Σ gives the set of expressions E.

[53]For ease of notation, we represent the RHS in 3.4 as a tuple, so that its components can still be edited within the subsequent pragmatic process (see below).

context (i.e. c) is needed in order to determine which of the here enumerated possible interpretations applies to a specific utterance.

Pagin and Pelletier claim that their compositional justification of moderate contextualism conflicts with both radical contextualism (as is evident from the compositionality requirement to which they adhere) and minimalism as defined by Cappelen and Lepore. For the latter have maintained that there is no such thing as a (systematic) theory of speech act content[54], which conflicts with the compositional pragmatic function.

The problem with the compositional pragmatic function, is that it does not give any reasons as to why it is likely that a function $\mathcal{M}$ exists that can make the relevant modulations, conditional on elements to be drawn from c. For the compositionality criterion does not explain *how* the context influences meaning. The modulation performed is a mere symbol, we take it to map context-independent meanings onto context-dependent meanings, but how the elements contained within c impinge on the choice over the precise mappings that $\mathcal{M}$ makes, remains unspecified. Therefore, if Pagin and Pelletier want to oppose Cappelen and Lepore's pessimism with respect to theoretic-/formal accounts of speech act content, then they must first show that both c and $\mathcal{M}$ can, at least for some cases, be construed along theoretically laid down lines. Moreover, they will have to show that there is a finite n-tuple that characterizes the context of utterance in such a way that the context-dependent meanings of all sentences come out right.[55]

As long as these issues are not addressed, minimalists will not be impressed by Pagin and Pelletier's strand of moderate contextualism. Radical contextualists will not be convinced either, for they will simply maintain that although the formulae do indeed characterize compositionality, examples like (3.19) make it clear that – at least in some cases – a non-compositional modulation function is mandatory in order to arrive at the right meaning.

3.4.1 Defining positions in terms of compositionality

In the light of the above discussion, I think the oppositions that play a role within the minimalism-/contextualism-debate can best be described in terms of the extent to which they adhere to the compositionality criterion. The two schemata that were given by Pagin and Pelletier, introduced in Section 3.4, allow me to formulate the programmatic stance that I will term *truth-conditional contextualism*:

- **Truth-theoretic contextualism** marks the hope that it will, in the end, be possible to give a compositional characterization of both the semantic composition function and (at least a substantive part of) the modulation function.
- **Semantic minimalism** is the claim that this is only possible for (a very narrow implementation of) the semantic composition function.
- **Radical contextualism** is the claim that not even the semantic process can be defined by a compositional function.

[54] See [10, p. 190] and [10, p. 200-1]; and its treatment in [35, p. 36].

[55] Needless to say, I do not think this is a viable task. See my critique of the similar contextual n-tuple that Kaplan uses in Section 3.2.3.

Truth-theoretic contextualism, according to the above characterization, is clearly the most ambitious of the three stances. It encompasses a larger area of linguistic behaviour than minimalism does, and assigns a much bigger role to formal approaches in the analysis of meaning than radical contextualism does.

But so far we have only made a characterization of the various stances. We have not yet shown that a truth-theoretic contextualism, i.e. a formal approach that explains context-sensitivity in language, is indeed possible. This is done in the next chapter.

CHAPTER 4

DISSOLVING THE CONTEXTUALISM-MINIMALISM DEBATE

4.1 Overview

In this chapter I give the solution for the minimalism/contextualism-debate that we described in Chapter 3. We do this by making use of the framework of truth-theoretic semantics under the methodology of radical interpretation (that was described in Chapter 2). The solution will adhere to the compositionality criterion that was introduced in Section 3.4. The solution consists of an alteration of traditional truth-theoretic semantics, so that it accounts for the problematic exemplars that radical contextualists have come up with, while adhering to the methodological rigour of the formal approaches that are characteristic of semantic minimalism (and some strands of moderate contextualism).

A first hurdle in providing the solution in terms of truth-theoretic semantics, is that although most minimalists (and a number of moderate contextualists) are already in the habit of employing this methodology in order to analyse meaning, such formal approaches are met with distrust by most radical contextualists. In Section 4.2 I will show that such reservations are due to a misconceptualization of what truth-theoretic semantics in actuality consists of. The section is kept brief due to spatial limitations but does cut some corners in destabilizing the radical contextualist's resentment against the method, thus paving the way for the utilization of truth-theoretic semantics in the rest of the chapter.

In Section 3.3.2 we saw that the notion of propositionality has, needlessly, been left inexact within the minimalism/contextualism-debate. At the same time, the notion seems central to formal theorizing in semantics, and so I will devote Section 4.3 to explaining the notion of propositionality as it figures in my theory of truth-theoretic contextualism.

In Section 4.4 I explain why radical contextualists seem to require so much more from a theory of meaning than minimalists do. The reason for this is that radical contextualists want to lay bare the fine-structure of many conceptual distinctions that do not contribute to communication at all, thus putting too much requirements on a theory of meaning. Needless to say, truth-conditional contextualism will not seek to uncover every conceptual distinction in the universe, but only those that are relevant with respect to a concrete instance of communication.

In Section 4.5 I give the solution for the background assumption problem that was introduced in Section 3.3.3 (it is actually a partial solution to the more general problem of the contextual dependency of meaning). I will show that the problem does in fact not occur for sentences that

are in active use. The problem does occur for those expressions that are not in active use, but do belong to the language in the formal sense (i.e. they are generated based on the same syntactic characterization). But this is only natural once we recognize that use is a necessary condition for interpretability.

Continuing the approach espoused in Section 4.5, in Section 4.6 I introduce the notion of a recurrent use, as well as that of a practice. Both contextalists and minimalists do not incorporate the way in which meaning is embedded within a broader context of use, yet this is the way in which the problem of context dependency can be resolved. In Section 4.6.2, the newly introduced concepts are applied to a concrete example of linguistic use that involves contextual supplementation of meaning.

In Section 4.7 I show how the considerations that occur in Section 4.6.2 must be incorporated into truth-theoretic semantics generally. The alteration of the theory will, conveniently, leave most of the usual formalizations that comprise truth-theoretic semantics intact. But although the reconceptualization I propose has little effect on the formalism, it will involve a fundamental revision of the way in which a truth predicate is to be characterized.

From the way in which the aspect of linguistic behaviour gets incorporated into the theory it will be clear that the formal approach to meaning is the right one, but that the way in which the formal theory must be related to communicative practices is very different from the way in which this has been traditionally conceived.

4.2 Removing radical contextualist's scruples towards truth-theoretic semantics

Radical contextualists often have reservations towards truth-theoretic semantics. This is not surprising, since they often repudiate the existence of a compositional meaning theory, of which truth-theoretic semantics is an instance. The objections that radical contextualists have towards truth-theoretic semantics are hard to make sense of since they so often seem to misrepresent the standpoint. We shall illustrate this with the following quote by Recanati:

> [...] consider the adjective 'red'. Vagueness notwithstanding, it expresses a definite property: the property of being red or having the colour red. That property could, in principle, go into the interpretation of a sentence in which the adjective 'red' occurs. (For example: 'Imagine a red surface.') But in most cases the following question will arise: what is it for the thing talked about to count as having that colour? Unless that question is answered, the utterance ascribing redness to the thing talked about (John's car, say) will not be truth-evaluable. It is not enough to know the colour that is in question (red) and the thing to which that colour is ascribed (John's car). To fix the utterance's truth-conditions, we need to know something more – something which the meanings of the words do not and cannot give us: we need to know what it is for that thing (or for that sort of thing) to count as being that colour. What is it for a car, a bird, a house, a pen, or a pair of shoes to count as red? To answer such questions, we need to appeal to background assumptions and world knowledge. [43, p. 183]

There are quite some mistakes in this passage. Recanati says that the word 'red' describes a 'definite property'. The definitite property of 'red' is said to be 'the property of being red or having

the colour red'. Now this is interesting, since Recanati is here (almost) using a disquotational definition. This ought to automatically answer the question he subsequently poses: 'what is it for a thing talked about to count as having that colour?' Recanati already gave the answer: what it is for a thing talked about to count as red is what it is for a thing talked about to be assigned the property of being red or having the colour red. This trivially fixes the truth-conditions of sentences of the form 'x is red' (provided object- and meta-language are the same).

Of course there will be circumstances in which we are unsure whether to describe a certain object as being red or not, but such considerations regarding the vagueness of the object language expression 'x is red' will be matched by the vagueness of the meta-linguistic property ascription of being red or having the colour red (again, provided object- and meta-language are the same).

Then there is the sentence: "It is not enough to know the colour that is in question (red) and the thing to which that colour is ascribed (John's car)." I don't see how one can know a colour or a thing. We only have knowledge of propositions, and a proposition is neither a property nor an object, but a combination of the two (in which the former is predicated of the later). Recanati's problem only crops up because he does not recognize the priority of the sentence in the process of assigning meaning (and this is a failure that is common to many radical contextualists, see Section 5.2.1).

First Recanati wants to know a colour *not applied to any object* (e.g. red), and he wants to know an object *irrespective of any properties* (e.g. a car). Only subsequently does he want to know the application conditions of the former with respect to the latter. But there is no such thing as a colour that is not applied to an object, neither is there anything like an object bereft of any properties. This way of conceptualizing semantic structure has been called the building-block theory of meaning. It has been expressly criticised by many, and truth-theoretic semantics was (amongst others) meant to provide a viable alternative to it. (For a characterization of what is wrong, from the viewpoint of radical interpretation, with a theory of meaning that starts out with words, see Section 5.2.1.)

From the reality of linguistic communication it is evident that talk of coloured objects takes priority, and that talk of properties and objects, taken in isolation, occurs only in abstracting over multiple instances of communication in which the isolated constituents always occur together. Once the priority of the sentence in the analysis of meaning is fully appreciated, and the language in which the RHS is framed is remembered to be strictly stronger than the language in which the LHS is framed, contextualist's scruples towards truth-theoretic semantics will dissolve.

4.3 Propositionality

Since the matter of propositionality has been so problematic in the past, see Section 3.3.2, I will expand on this in the light of my account of truth-theoretic contextualism. This section's main intent is to subvert the various minimalist and contextualist misconceptions.

Propositions are those entities that function as the meanings of sentential expressions. In truth-theoretic semantics, as we have described it in Chapter 2, we do not ponder over the ontological reality of meanings, and so neither do we ponder over that of propositions either.[1] But nothing turns on this, since the same problem recurs in the purported inability of meta-language expressions to interpret object language sentences. For instance, if 'Steel isn't strong enough' would need the

[1] From this it follows that in truth-theoretic semantics we cannot appeal to propositions as the locus of e.g. opaque contexts, see [15].

added phrase 'for carrying a construction of type X' in order to be propositional, then this would, translated in terms of truth-theoretic semantics, mean that the T-sentence (4.1) does not suffice for interpretation. The insufficiency of the T-sentence would first have to be shown by the theory's being confronted with negative instances of empirical validation.

'Steel isn't strong enough' is $T(S, t)$ iff steel is not strong enough at t. (4.1)

The word 'proposition', as I use it, therefore denotes the meta-linguistic expression that occurs within the RHS of a T-sentence. As long as the set of sentences is, as we have been assuming all along, defined by a finite lexicon and a finite number of recursive rules for grammar, the set of 'propositions' is defined – accordingly – through the T-sentences that our truth theory generates. What set of meta-linguistic expressions are labelled 'propositions', in the here described sense, is therefore determined by the formal syntactic and semantic specification of the language that we are interpreting.[2]

Since there are many ways in which we can formalize (a certain aspect of) a natural language, there will be many different sets of propositions associated with a single natural language (through the vehicle of different formal languages that abstract away from their natural counterparts in different ways). Not only does the proposition that gets associated with a specific natural language sentence differ between disparate formal abstractions of the same natural language; but also whether or not there exists a proposition that corresponds to a certain collection of natural language signs, is something that can only be determined with respect to the specific way in which the natural language is being formalized (because there may be multiple ways in which natural language sentences get paired with formal items of type 'truth-value').

In Section 3.3.2 we saw that incompleteness intuitions focus exclusively on intuitions regarding natural language expressions (such as (3.13)). But this, in the light of the relativity of the set of propositions that was sketched above, is a wrong way of conceptualizing the matter. Whenever two researchers disagree over whether a natural language sentence S expresses a proposition or not, they might as well have different formal analyses of S in mind (i.e. different definitions for sentencehood).[3]

Let us consider an example in which different propositions are assigned to the same sentence. We already saw that the sentence 'It's raining' is taken to express disparate propositions according to different authors (see Section 3.2.2). Minimalists think of something like (4.2)[4], Borg [6, p. 230]'s liberal truth-conditions (adding an extensionally quantified argument place) would make us come up with something like (4.3). Moderate contextualists may come up with something like (4.4) or (4.5) (see e.g. Stanley [50, p. 240]). And then there are the perspectival contextualists (discussed in Section 3.2.5) who hold that no proposition can ever be complete enough so as to describe an utterance's meaning in context (maybe only an proposition that describes a complete world-state would suffice in their opinion).

[2]From these conditions it follows that the language with respect to which the set of propositions is defined can never be a natural language, but is always a formal abstraction of such a natural language.

[3]It must be noted that even though what is a proposition is now seen to depend upon the formalization that is chosen, the characterization of this set is itself still stringently defined. A formal language that is agreed to characterize both syntax and semantics of natural language expressions determines (in a stringent manner) the total set of propositions.

[4]See Cappelen and Lepore [10, p. 66].

That it is raining. (4.2)

That it is raining somewhere. (4.3)

That it is raining in Amsterdam. (4.4)

That it is raining in the Derkinderenstraat, on house number 101's balcony. (4.5)

Given the necessary condition on a truth-theoretic semantics that the meta-language must be strictly stronger than the object language, when I encounter someone who speaks my language, I will use (4.2) in order to interpret him. So far my choice seems to echo Cappelen and Lepore's analysis, but observe that I did so under the provision that the person I encounter *speaks the same language as I do*. But this is not found out by merely looking at the signs he utters, since whether or not someone speaks the same language as I do, *can only be found out in interpretation*.

Elaborating on this line of thought, if the meta-language expression 'It is raining' does not seem to match the speaker's use of the object language sentence 'It's raining', i.e. if the speaker's language differs from the interpreter's language with respect to the word 'rain', then additional aspects of the situation in which the speaker is uttering those words must be incorporated into the T-sentence's RHS in order for communication to succeed.

Communication is successful as long as agreement between speaker and interpreter is maximized. So if the meta-language expression 'It is raining near S' makes speaker and hearer agree more often than the minimalist's expression 'It is raining' does, then the RHS is legitimately extended beyond the disquoted version.

Since there is no way for the interpreter to anticipate in advance all of the possible ways in which the speaker's language might differ from hers (even with respect to a single word like 'rain' this is already the case), it is simply not sensible to ask for the proposition expressed by a sentence (like 'It's raining') *once and for all*, i.e. *irrespective of use*. Yet both contextualists and minimalists seem to believe it is possible to do so. They saliently take the language for which they are defining the propositional contents of sentences to be the English language. But the English language is not a static entity of the sort they make it out to be. For not only do grammar and lexicon change intermittently (between speakers, between situations, and over time), also the use to which a static grammar and a static lexicon might be put is itself not static.[5]

Summarizing, the RHS of the T-sentence for 'It's raining' can be any of (4.2) through (4.5). The specifically chosen RHS-description is conditional on the speaker and hearer that are involved in the act of interpretation, relative to a certain communicative purpose, embedded within a practice in which both agents figure. Definitorally speaking, *a non-propositional content is an RHSs expression that is suboptimal with respect to communication*.[6]

[5]Under some conditions conventions can be made to restrict linguistic use of course, but the philosophically relevant point is that such conventions are neither necessary nor sufficient for pinpointing the meaning with respect to a static syntactic structure. This topic is treated in Section 5.4.

[6]The shift from defining things in terms of truth to defining the same things in terms of the interpretative potential of a truth theory, was already observed in Section 2.5.

4.4 The depth of analysis that is necessary for uncovering meaning

In Section 3.4.1 we saw that truth-theoretic contextualism was characterized as the hope for the invention of a future account of the compositional functions that give the meanings of sentences. But radical contextualists have oftentimes provided examples of sentences that, they claim, cannot be represented by *any* compositional modulation function. Such counterexamples are instantiations of schema (4.6).

$$\mathcal{M}(\langle \rho(\sigma), \mu(t_1, c), \ldots, \mu(t_n, c)\rangle) \neq \langle \mathcal{M}(\rho(\sigma)), \mathcal{M}(\mu(t_1, c)), \ldots, \mathcal{M}(\mu(t_n))\rangle \tag{4.6}$$

One such an example is (4.7), given by Recanati [43, p. 178]. According to Recanati this example shows that non-compositional modulations apply directly to the contents of the constituents of sentences. Modulatory pragmatic intrusion enters into the semantics of (4.7) because "the process of representational transfer must take place *before* the composition rule associated with the noun-noun construction applies to the semantic values of the nouns 'stone' and 'lion'."[7]

There is a stone lion in the courtyard. (4.7)

Using the compositional functions introduced in Section 3.4, we get (4.8). In this formula $\rho(\sigma)$ represents the composition function that applies (from top to bottom in the parse tree) ‖there is a stone lion‖ to ‖in the courtyard‖, ‖there is a‖ to ‖stone lion‖, ‖stone‖ to ‖lion‖, and ‖in‖ to ‖the courtyard‖.

$$\mu(\sigma(\text{there is a, stone, lion, in, the courtyard})) = \tag{4.8}$$
$$\langle \rho(\sigma),\ \lambda P.\lambda Q.\exists x(P(x) \wedge Q(x)),\ \lambda P.\lambda x.Stone(P(x)),\ Lion,\ \lambda y.\lambda x.In(x, y),\ c_0\rangle$$

Recanati notes that "the thing that is said to be in the courtyard is not a (real) lion but a *representation* (more specifically, a statue) *of* [a] lion. What is said [in (4.7)] to be made of stone here? Clearly, it is the statue, rather than the lion which the statue represents."[8] So what Recanati has in mind as the result of the representational transfer involved is the added qualification that the stone lion is not a real lion made of stone, but a real lion's stone statue. This instance of representational transfer takes us from (4.9) to (4.10) as the meaning of (4.7).

$$\exists x(Stone(Lion(x)) \wedge In(x, c_0)) \tag{4.9}$$
$$\exists x, y(Stone(Statue(x)) \wedge RepresentationOf(x, y) \wedge Lion(y) \wedge In(x, c_0)) \tag{4.10}$$

But is this really a counterexample to the viability of the compositionality principle in the explication of meaning? If we employ Pagin and Pelletier's conceptualization of the matter it clearly is, for (4.10) cannot be composed out of the meanings that are attributed to the subsentential

[7] See [43, p. 178].

[8] [43, p. 178].

components involved.[9] But I want to call to mind the distinction, introduced in Section 2.3.1, between logical structure and the analysis of concepts.

The idea is that if we keep in mind to what depth the excavation of grammar ought to be performed in order to uncover the skeleton of logical structure that is needed in order to model successful communication, we may be able to rebut purported counterexamples such as Recanati's.

The problem with the representation that is used in (4.8) is that it is geared too much towards conceptual analysis. But maybe we need not dig so deep in order to get to the level of semantic perspicuity that a theory of interpretation requires. All we need, from the point of view of truth-theoretic semantics, is an axiomatic specification of the satisfaction relation that generates, among others, the T-sentence (4.11).

'There is a stone lion in the courtyard' is $T(S,t)$ iff there is a stone lion in the courtyard near S at t. (4.11)

It is very unlikely that in order to generate (4.11) we need to distinguish between stone (living) lions and stone statues representing a lion. As long as names for object language sentences get paired with corresponding meta-language sentences, all according to a finite scheme, there is – as far as meaning is concerned – nothing missing.

Another example will convey the same point even more clearly. Suppose that at some point in an interpreter's learning the English language, in the truth theory that models her linguistic competence, the word 'and' gets interpreted on a par with the logical connective for conjunction. We assume that her theory already assigns the correct meanings to sentences (4.12) and (4.13) (whose meanings we will take to be their mere disquotations in the meta-language, but this is not essential to the argument).

John and Mary had a child. (4.12)

John and Mary got married. (4.13)

In addition, we assume that her theory can handle basic forms of NP-ellipsis. Now, according to her theory, the meaning of (4.14) is exactly the same as that of (4.15). This is only natural given her unfortunate analysis of the word 'and'.

John and Mary had a child and got married. (4.14)

John and Mary got married and had a child. (4.15)

It is evident however that utterances of (4.14) and (4.15) have quite different speech act contents, and that the mere conjunction is not what the English word 'and', generally speaking, brings across in communication. If we are only interested in communication, i.e. in the way in which successful reciprocal interpretation proceeds, we can simply replace the meta-language sign of conjunction with the word 'and' (provided this word is part of the interpreter's vocabulary too).

[9]The problematic aspects of Pagin and Pelletier's approach were discussed in Section 3.4.

If we are making a thorough conceptual analysis of the word 'and', as it occurs in the English language, then we are not allowed to use the meta-linguistic 'and' in order to characterize the sequentiality that is expressed in sentences containing that word. (The reason why the word 'and' does not suffice from the point of view of thorough conceptual analysis, is that the word 'and' is no part of any logic, supposing that it is a logic in which the analysis is to be carried out.)[10]

So whenever a radical contextualist comes up with an example that seems to contradict a compositionally structured meaning theory, we ought to first ponder the question whether the fine structure the example hinges on is indeed essential for the task of axiomatizing an interpretative truth theory. In order for an example to be a counterexample of the sort intended, its consequences would have to range over the potential infinitude of possible truth-theories (needless to say, counterexamples of such intricacy have not yet been given).[11]

4.5 Dissolving the background knowledge-problem

Radical contextualists claim that T-sentences like (4.22) do not suffice to give the truth-conditions for the object language sentences that are named in the LHS. The idea is that, as far as meaning is concerned, something essential is still missing in the corresponding RHS expression.

'This car is red' is $T(S,t)$ iff the car indicated by S at t is red at t. (4.22)

[10]The structure that must be recognized for a theory that *interprets* (4.14) and (4.15) is provided by (4.16) - (4.18). The structure that must be recognized for a theory that *analyses* (i.e. seeks conceptual understanding) (4.14) and (4.15) is provided by (4.19) - (4.21) (where I have assumed the meta-language to be (first-order) predicate logic).

'NP_1VP_1 and VP_2' is $T(S,t)$ iff '$NP_1\ VP_1$' is $T(S,t)$ and '$NP_2\ VP_2$' is $T(S,t)$ and 'NP_1' and 'NP_2' co-refer. (4.16)

'John and Mary got married and had a child' is $T(S,t)$ iff 'John and Mary got married' is $T(S,t)$ and 'John and Mary had a child' is $T(S,t)$. (4.17)

'John and Mary got married and had a child' is $T(S,t)$ iff John and Mary got married prior to t and John and Mary had a child prior to t. (4.18)

'$NP_1\ VP_1$ and VP_2' is $T(S,t)$ iff '$NP_1\ VP_1$' is $T(S,t)$, and '$NP_2\ VP_2$' is $T(S,t)$, and 'NP_1' and 'NP_2' co-refer, and the event denoted by '$NP_1\ VP_1$' occurred prior to the event denoted by '$NP_2\ VP_2$'. (4.19)

'John and Mary got married and had a child' is $T(S,t)$ iff 'John and Mary got married' is $T(S,t)$ $\wedge$ 'John and Mary had a child' is $T(S,t)$ $\wedge$ the event denoted by 'John and Mary got married' occurred prior to the event denoted by 'John and Mary had a child'. (4.20)

'John and Mary got married and had a child' is $T(S,t)$ iff $ISA(e, marry)\ \wedge\ Subject(e, \langle John, Mary\rangle)$ $\wedge IntervalOf(e,i)\ \wedge\ i < t\ \wedge\ ISA(e', have_child)\ \wedge\ Subject(e', \langle John, Mary\rangle)$ $\wedge\ IntervalOf(e', i')\ \wedge\ i' < t\ \wedge\ i < i'$. (4.21)

[11]Introducing the Davidsonian distinction between logical structure and the analysis of concepts in the discussion over the viability of the compositionality criterion within the domain of meaning assignment has an application far beyond the present discussion of radical contextualism. As a matter of fact, the here described critique applies to a huge number of attempts to disprove "the Gricean idea that pragmatic processes operate globally on the output of the grammar" [43, p. 179]. Examples of such objections that ought to be revised in the light of what I have stated in this section include Sag [46] and Jackendoff [24, p. 65-6], but there are many others.

Radical contextualists think that such T-sentences, i.e. ones that exemplify pure disquotation (except for being supplemented for expressions from a quantitatively delimited set of expressions, e.g. the basic set), do not suffice in order to convey the truth-conditions of object language sentences. This is so because they do not understand that the *complete* T-sentence is what gets empirically validated in the process of radical interpretation.

The empirical validation procedure, rightly understood, applies to T-sentences as a whole and not to the RHS in isolation. In this way we are sure that the elements that are left unspecified in the object language sentence get paralleled by meta-linguistic descriptions that make use of expressions that convey the same forms of ambiguity and that exhibit the same level of vagueness.[12]

There are three ways in which (aspects of) expressions in the LHS can be handled in T-sentences. First, if object- and meta-language are the same (with respect to the sentence under consideration), then the elements that were implicit in the object language sentence will be implicit in the meta-linguistic description too.[13] An instance of such a sentence is 'Bordot is good' (i.e. (2.6) from Section 2.3.1).

The second way in which object- and meta-language can be related, is when they are different (with respect to the sentence under consideration). Under such conditions the RHS cannot be a simple disquotation (supplemented for expressions belonging to a qualitatively delimited set) of the sentences occurring in the LHS. The RHS description must either involve a supplementation of the disquotation, so as to make some of the missing pieces explicit, or it must be an altogether different sentence. An example of this is (4.23).

'Schnee ist weiss' is-T-in-German iff snow is white. (4.23)

The third (and last) possibility is one in which the object language sentence carries some implicit information that is not relevant for successful communicative interaction between speaker and interpreter to take place. This is the variant that was extensively discussed in Section 4.4. Examples were Recanati's stone lion, and my example involving the temporal sequentiality of 'and'.[14]

In Section 4.2 we saw that Recanati wants the meaning of 'That car is red' to be amplified by a specification of 'what it is for a car to be red'. But it is unclear why this would be the amplification that is needed. Clearly, we can ponder an infinity of questions, e.g. 'what it is for that car to be red on a Sunday afternoon?', etc.

Anticipating all the amplification-requests that a radical contextualist may come up with, the only thing that would satisfy the search for the truth-conditions of a sentence like 'That car is red'

[12]It is clear that for every natural language we can come up with a corresponding meta-language. Because we can always use the object language itself, to which predicate T is added. The challenge is whether we are able to recursively structure the meta-language to as to make a finite mechanism bestow meaning upon a potential infinity of expressions.

[13]Moreover, they are implicit in the RHS for the very same reasons as they are implicit in the LHS. So if the absence of a certain qualification in the RHS seems mysterious, then it is just the kind of absence (and just the kind of mystery) the LHS expression requires.

[14]It is not (always) possible to say for a given T-sentence which of the three possibilities here enumerated is at play. For this partly depends upon what implicit aspects the beholder takes to be part of the object language sentence. E.g. if one thinks that the German word 'Schnee' has associated with it some background knowledge that is not carried over by the English word 'snow', yet one does not think that the incorporation of this background knowledge is necessary for communication to proceed, then (4.23) is not only an instance of the second, but also of the third type.

is an exhaustive specification for each and every possible (in the sense of 'imaginable') instance in which the sentence could be legitimately applied. Apart from the ridiculousness (and impossibility) of describing everything there is within the universe at a given moment in time just in order to know the truth-conditions of the sentence 'This car is red', in carrying through this line of thought we have come to equate truth-conditions with truth-values (for we can now say for each and every state of the world whether the sentence 'The car is red' is true in it or not).

These mistakes all hang together with a persistent misunderstanding of what truth-conditions for a sentence ought to look like. The mistake is apparent from the following quote, in which the minimalist's viewpoint (drawing on Searle's example that was already discussed in Section 3.3.3) is attacked:

> We may know the obedience-conditions of 'Cut the sun' in a purely 'disquotational' manner (i.e. we may know that 'Cut the sun' is obeyed iff the addressee cuts the sun), without knowing what counts as cutting the sun, in the context at hand. [43, p. 185]

Firstly, it is not evident that the T-sentence for the object language sentence 'Cut the sun' would only consist of the disquotational version of that sentence (supplemented for context sensitive expression belonging to the basic set). This will only be the case if speaker and interpreter share the same language with respect to this sentence (i.e. the first of the above identified possibilities). If this is not the case, then alternative descriptions may be possible as well (i.e. the second and the third possibilities in the above). An instance of this is (4.24).

'Cut the sun' is obeyed by S at t iff S cuts the cardboard model of the sun at t. (4.24)

I will call sentences, like 'Cut the sun', that can be formed by the recursive syntax (and are assigned a meaning by the compositional semantic theory) but that have (as of yet) no concrete use in instances of communication, *bereaved sentences*. The reason why we do not know what counts as cutting the sun, beyond the meta-linguistic description that it will involve *cutting the sun*, is that we are not acquainted with any practice in which the sentence 'Cut the sun' has a viable role to play. But this is not remarkable at all, as long as we keep in mind that linguistic expressions receive their meanings in virtue of the practices in which they occur.

This is not to say that bereaved sentences have no meaning whatsoever. But the meaning that they do have is not correlated with use, so that the corresponding T-sentences are purely hypothetical, i.e. mere theoretic presuppositions. This is only evident once we recognize that in a meaning theory active language use bestows meaning onto potential language use. The axiomatization itself does not distinguish between bereaved sentences and those that are in active use. But since the procedure of empirical validation is necessarily restricted to only those sentences that are in active use, the meanings that the bereaved sentences receive is parasitical upon the meanings that are assigned to sentences that are in active use.

If one is still persuaded by Searle's arguments, then one must be cognizant of the fact that the same reasoning ought to be levelled against *every* bereaved sentence in the language. Since every grammatically correct nonsense sentence is a bereaved sentence too, e.g. the well-known sample (4.25), it is manifest what paucity of meaning must necessarily be assigned to sentences that do not occur in correlated practices.

'Colourless ideas sleep furiously' is $T(S, t)$ iff colourless ideas sleep furiously at t. (4.25)

It is not a deficiency of semantic theorizing that we do not know how to establish whether a flock of colourless ideas are sleeping furiously or not. And neither ought Recanati or Searle wonder what it would be like to cut the sun, or to sun the cut, or to cut pondering an orphanage, etc.

This all follows from the way in which truth-theoretic semantics proceeds. The axioms define a satisfaction relation that makes (a large enough amount of) sentences that are empirically validated in discourse come out correct (through the vehicle of generated T-sentence hypotheses). The truth theory is validated with respect to a finite number of sentences, while the satisfaction relation (except for the most rudimentary of languages) will bestow meaning onto an infinite number of potentially spoken sentences. The T-sentences that are given for bereaved object language sentences are theoretically indistinguishable from those that are given for object language sentences that are in active use. But from the point of view of communication, the former have the character of *unproven hypotheses*, whereas the latter's correctness has already been verified.

Most sentences will, in actuality, never get used and will receive 'strange meanings'. The strangeness of meaning corresponds to the unlikeness of uttering the sentence within a specific discourse. As soon as a new sentence gets uttered in a discourse, and the hypothetical T-sentence that the axiomatization generates gets invalidated by empirical evidence, then the truth theory changes in order to make the corresponding T-sentence pass through the empirical validation process that consists of the newly invented use. The sentence, upon use, will lose its strangeness of meaning accordingly, as was indicated by sample sentence (4.24).

Therefore, there is nothing peculiarly odd about the oddness of meanings for unused or bereaved sentences. As long as oddness gets transformed into clarity upon use, this is just the behaviour our theory ought to describe.

4.6 Practices

Remember Bezuidenhout's apple sorting example, introduced in Section 3.3.1, from which it became apparent that there is really no end to the number of ways in which an apple can be (said to be) red. As a consequence of this, there is an indefinite number of (different) truth-conditions for the sentence 'The apple is red'. Observations like these, generalizable in order to apply to the whole of language, were seen to depend upon our being able to *imagine* situations in which a sentence could be used in a variety of ways (e.g. in Bezuidenhout's example these are imaginary apple sorting tasks). It is therefore *imagination* that Cappelen and Lepore focuss on in their formulation of the principle according to which moderate contextualism collapses into radical contextualism:

> With sufficient ingenuity a context shifting argument can be provided for every sentence.[15]

But even though there is such a pluriformity of ways in which an apple can be imagined to be red, it is an altogether different matter whether in the language that we actually use, we would not describe (at least the vast majority of) all these cases by saying – simply – *that the apple is red*, i.e. (4.26).

[15] Taken from [10, p. 40].

'The apple is red' is $T(S,t)$ iff the apple indicated by S at t is red at t. (4.26)

The RHS in (4.26) could potentially have been a whole lot longer than this, incorporating all kinds of extra conjuncts, corresponding to some of the fancy distinctions that could be explicitly uncovered by any number of imaginary examples (e.g. the percentage of the apple's area that is red, the lightning conditions under which the observation is made, etc). But as long as the incorporation of these extra qualifications (of an apple's being red) do not increase the evidential base that is covered by our truth theory, we do not feel the need to extend the RHS of (4.26).

That under some imaginary use I should have employed a different T-sentence in order to interpret the same object language sentence does not bother me as long as I am concerned with active linguistic use solely. In the examples that are given by Travis, Bezuidenhout, and many others, we are confronted with new (and sometimes unusual) contexts. We are then asked whether the meanings that our current truth-theory gives us will also apply in these imaginative (or non-actual) situations. But what is not recognized, is that it is only evident that our current theory does not issue in T-sentences that will (always) work in the newly invented situation. The meanings that a truth theory generates depends upon the empirical instances with which the theory has in the past been confronted. This means that different interpreters employ different truth predicates while speaking the same language. The distinction depends on what are the sentences that have been actively used in an interpreter's past communicative experience.

The problem was extensively discussed in Section 4.5, namely that in imagined practices now-bereaved sentences may become active.[16] Under different uses the empirical verification set will differ too. It follows that the axiomatic specification of the truth predicate will vary accordingly.

The main problem of the minimalism/contextualism-debate can now we characterized as a disregard for the philosophical importance of the inherent structure of the (imagined) situations in which sentences get uttered. I shall therefore start by fleshing out the as of yet unarticulated role of (real or imagined) situations in the process of uncovering meaning.

4.6.1 Practices and recurrent uses

Radical contextualists argue that there is an indefinite number of distinct uses to which a single expression might be put. In arguing this, they draw on the intuition that a single sentence can play an indefinite number of distinct roles in different (partially) linguistic uses or 'practices'.

If a T-sentence can be construed by making use of information that is available to the interpreter within the context of utterance (or within the set of utterance-contexts that belong to a discourse), there must be what I call a *recurrent use*. A recurrent use is always correlated with an object language sentence (or with a tiny object language discourse). It is a teleologically structured conglomerate of (partially) linguistic behaviour that the interpreter thinks of as central to the specific use to which an utterance of the object language sentence under scrutiny is being put.

Recurrent uses can therefore be thought of as packages of behaviour descriptions that an interpreter takes to be correlated with utterances of certain object language sentences (or with utterances

[16] Even if the same set of sentences stays active, a change in use potentially involves a change in the meanings that are assigned to this stable set of sentences.

of tiny object language discourses). We need not dig too deep into the foreign here, descriptions of recurrent uses are the kind of things that typically end up inhabiting the RHSs of T-sentences.[17]

Firstly, it ought to be pointed out that the above description of a recurrent use does not imply that an interpreter must grasp a certain use prior to her getting involved in instances of communicative interaction. It is clear that learning to use a new sentence (or an old one in a new way) largely consists of coming to understand the teleological fine-structure that binds otherwise isolated instances of behaviour, together with contextual aspects, into a coherent package. Therefore, the understanding that an interpreter has of a linguistic interchange is always relative to the extent in which she has come to grasp the teleological fine structure of the linguistic behaviour conglomerate in which she finds herself immersed.

Neither does the above imply that every recurrent use must have a clear goal and/or a clear procedure leading up to that goal. Small talk is an instance of linguistic interaction that is structured to a certain extent (and therefore recognizable), while lacking a definite goal. It should not come as a surprise that some recurrent uses have clear goals and accomplishment-lines leading up to their goals, whereas others do not.[18]

Lastly, it is important to emphasize that recurrent uses do not have an existence outside of the interpreter's mind (only behaviour does). A recurrent use is what an interpreter has to inventively come up with all by herself. Neither is it necessary that the recurrent uses that the interpreter invents are identical to those put forward by the speaker (or the community of speakers). This is recognized once we are reminded of the fact that truth theory coincidence, as far as everyday communication is concerned, can only be determined at the level of T-sentences (see Section 2.4).[19]

The notion of a recurrent use points to the ability of the interpreter to group various forms of observed behaviour into purposeful packages, and the notion of a practice then designates what consequences certain decisions for packaging the material in specific ways has for the resulting truth theory. The notion of a practice sets in whenever T-sentence construction, based on the interpreter's recognition of recurrent uses, is unable to produce a consistent truth theory of acceptable complexity. This means that theoretical considerations are to do the job of practice individuation. (The specifics of this will be explained below.)

Because of the way in which the concept is defined here, practices cannot be available at the onset of interpretation. This means that practices are not to be looked upon as *ex machina*-concepts in terms of which we can describe what goes on in (radical) interpretation. Instead, the question of what constitutes a practice ought to be answered within the process of interpretation itself (or so I will argue).

[17] If a T-sentence gets invalidated by empirical evidence, the RHSs of certain T-sentences must be altered in the light of this. And one way of altering T-sentences is by attempting to find new conglomerates of behaviour, i.e. new recurrent uses that can come to replace the old RHS.

[18] E.g. unclarity regarding a recurrent use's goals can itself be a characteristic of that recurrent use.

[19] In discussing the interaction between a speech community and an interpreter, we touch on the fundamental topic of language normativity. If what the interpreter has to construe does not have to coincide with some preliminarily available norm residing in the speaker (or in the community to which the speaker belongs), then standard accounts of normativity and rule-following do not apply. The topic is not peculiar to our treatment of context dependency, but occurs in radical interpretation generally, and will therefore not be elaborated upon here.

4.6.2 Apple-sorting practices

Let us apply the newly instituted concepts from Section 4.6.1 to the apple-sorting task described by Bezuidenhout in Section 3.3.1. The example introduces two different practices in which the word 'red' gets used. But let us first observe the situation in which we are not concerned with a specific use. If there is no such thing as a concrete use that hinges on the distinction between an apple's being red on the outside as opposed to its being red on the inside, then there is no need to introduce T-sentences that specify anything but the ordinary, i.e. (4.27).

'The apple is red' is $T(S, t)$ iff the apple indicated by S at t is red at t. (4.27)

As long as we imagine no specific apple-sorting task, the successful completion of which is based on specific ways in which an apple may be (said to be) red, (4.27) suffices. Under normal conditions the fact that an apple is red on the inside as well as the fact that an apple is red on the outside, are both captured by the meta-linguistic phrase *that the apple is red*. Consequently, (4.27) will be validated by instances of empirical evidence like (4.28).

John holds true 'The apple is red' at t, and the apple is red at t. (4.28)

It is only when evidence is found that contradicts our T-sentence (4.27), that we are forced to include some further specification as to an apple's being red. This is where we turn to Bezuidenhout's example (though any example concerning distinct uses of the same expression will do of course). It is by imaginatively invoking two different recurrent uses, one in which apples are to be sorted with respect to their external redness and one in which they are to be sorted according to their internal redness, that the necessity of incorporating a further specification of an apple's redness within our T-sentences arises.

Under the imaginative red fungus that affects some apple's insides, there are apples that are red on the inside but not on the outside, and vice versa. Now it is evident that (4.27) will no longer be validated by arbitrary instances of (4.28). The reason for this is that there are apples for which the truth-conditions are orthogonal between the two tasks (namely those apples that are red on the inside but not on the outside, or the other way round).

One might object to this by maintaining that the meaning of the word 'red' in the meta-linguistic description in (4.27) already contextually depends upon the practice that is involved, and the RHS therefore matches the meaning of the object language sentence 'automatically'.

However, it is clear that this contextually dependent interpretation of the word 'red' does not conform to the requirements of radical interpretation. Even though a theory involving this alternative reading of the word 'red' would issue in the right truth-conditions, we must also keep in mind that a truth theory that is to do the job of a meaning theory must put *every* rational agent in a position to interpret any of a potential infinity of utterances. But it is clear from the discussion in Chapter 3 that the meta-linguistic word 'red' is not conditional on the practice involved in an introspective way, and yet this is what radical interpretation requires.

Minimalists, for instance, claim that the word 'red' is conditional on nothing at all, let alone a 'practice'. And they are justified in their sceptic stance; for even if one does have the intuition that

the meaning of the word 'red' depends on context, then one still has not provided an account of when and why either of the supplemented meanings ought to be invoked (at least nothing nearly as concrete as T-sentences).

In short, an account of the context dependency of meaning must show how the meanings of sentences depend on aspects of context in a way that does not rely on intuitions that some (claim to) have but others (claim to) lack. Instead, the resulting theory ought to put both minimalists and contextualists (as well as everybody else) in a position to interpret the word 'red' conditional on whatever intuitions one might have.[20]

If we keep fixed the apple-sorting task under consideration, it is quite easy to come up with the right T-sentences.[21] On the former task, 'This apple is red' is only legitimately applied to apples that are red on the outside, as in (4.29).

'The apple is red' is $T(S, t)$ iff the apple indicated by S at t is red on the outside at t. (4.29)

Now our apple sorter switches to the other task. Provided she transposes the T-sentence she constructed during the performance of the former task to the current endeavour, she will be likely to run into problems at the very onset of performing her task. For every apple whose outside does not reflect its inside, the truth conditions of 'The apple is red' will be precisely orthogonal to the circumstances in which the utterance in the present practice is deemed appropriate. The right T-sentence under the second task is of course (4.30).[22]

'The apple is red' is $T(S, t)$ iff the apple indicated by S at t is red on the inside at t. (4.30)

But in real interpretation we cannot presuppose the task the interpreter is involved in. This means that by introducing the possibility of using the same words in different task-contexts, the interpreter's burden is seen to increase. For in addition to finding the right axiomatic specification of the truth predicate that generates the right T-sentences, the interpreter must now also acknowledge – in some sense of the word – the task that she is involved in.

Sticking to our apple-sorting example, the interpreter in addition needs to understand the two practices that she is involved in, and she needs to understand that they are practices that differ from one another. Provided these identifications are made, it becomes theoretically introspective how an interpreter is able to switch between two practices at the drop of a hat without the need to postulate an indefinite number of contextual parameters.

Let us see how the above considerations apply to the apple sorting example. Suppose my working partner and I are given two tasks: he is to count the number of apples that are red on the inside,

[20] So in order to make communication succeed, a minimalist might need different T-sentences than a contextualist does, since their disparate intuitions already indicate that their meta-languages differ.

[21] Note that whether a T-sentence is 'right' or not does not involve an appeal to intuition, but is defined in terms of success in confronting instances of empirical validation (see Section 2.4).

[22] Even though this is the evident answer, it is clearly not the only way in which our interpreter could have altered her theory. Alternatively, she might have used a temporal variable in order to define a term 'red2' that characterizes the property of being red on the outside until 5 p.m., but red on the inside afterwards. Assuming she always switches between the two tasks, in the same order, at 5 p.m., this would seem a viable (though somewhat strained) possibility. (The problem with this interpretation is of course that it is susceptible to minor alterations in the tasks involved.)

whereas I am to count the number of apples that are red on the outside. We check each other's categorizations, by correcting the other as soon as he utters the sentence 'The apple is red' illegitimately. Suppose my working partner seeks some variation in his dreary task sometimes, so that whenever he feels like it, we switch tasks (so that he counts the apples that are red on the outside, and I count those that are red on the inside). I am still able to identify the task he is intending me to follow, since in one task he will be slicing apples open, whereas in the other task he will leave them intact. What truth conditions I apply in order to interpret his utterances is thus conditional on this rather easy-to-make identification.

Although circumstantial regularities are sometimes sufficient for identifying a practice (as above), such circumstantial regularities are never necessary. For instance, even if my working partner is quite a bit weirder than previously supposed, slicing apples open irrespective of which of the two tasks he intends us to perform, I should get the hang of what truth conditions to use quite easily. My utterances of the sentence 'The apple is red' will be observed to fail every time my working partner silently switches tasks, thus by invoking a negative reaction from him, I am able to switch truth theories based on a single reprimand.

The example clearly bears out the distinction between recurrent uses and practices. For whereas what is to count as a recurrent use is a subjective determination involving intuition, what is to count as a practice is something the truth theory that is employed sets forth (conditional on a preliminary subjective identification of recurrent uses of course).

It follows that the identification of practices depends on the peculiarities of the language user involved, namely it depends on what recurrent use the interpreter identifies, as well as on what instances of linguistic behaviour she has been confronted with.[23] The problem of the contextual dependency of meaning is thus solved in a theoretically introspective way (probing beyond mere intuition), without making the theory inert with respect to the peculiarities of the interpreter involved.

The above example already allows a preliminary critique to the basic tenets that characterize the minimalism/contextualism-debate. Cappelen and Lepore [10] despair that once we have started down the slippery slope of adding additional descriptions of the conditions of an apple's being red, we will soon be inclined to specify an indefinite number of such conditions regarding the redness of apples. But this is an illegitimate inflation of the problem. According to my treatment, appealing to an indefinite number of contextual factors would imply an indefinite number of uses in which disparate specifications of an apple's redness would have interpretatively distinctive purposes. But no speaker is involved in an indefinite number of practices at any one time!

4.7 Truth-theoretic contextualism

In Section 4.6 we saw how a truth theory could be made to work for context sensitive sentences in the face of the use to which these sentences were being put. In this section we are concerned with

[23] An example will make the interdependence of the subjective recognition of recurrent uses and the objective identification of practices even more clear. Though warning a person to not get too close to a dog with sharp teeth and bad temper on Friday and warning a person for the same thing on Tuesday, may be called different *recurrent uses*, it is unclear how these could be different *practices* (unless the interpreter has a theory that is very different from mine to begin with).

the theoretical revisions that are necessary in order to make the approach that was only illustrated with respect to a single example, work for all of linguistic behaviour.

Central to the example in Section 4.6.2 was the identification of a practice. In sticking with the original formulation of truth-theoretic semantics, the fact that the truth-conditions for an object language sentence can vary according to the practice in which that sentence is being used, may be thought to be implemented by adding terms that denote practices within the RHSs of the T-sentences the theory generates. Within the thus altered scheme, the RHSs of T-sentences can be thought of as conjunctions of conditional statements that have identification conditions for practices within their antecedents, as e.g. in (4.31).

'The apple is red' is $T(S, t)$ iff the apple indicated by S at t is red on the outside at t, in case of P_1; or the apple indicated by S at t is red on the inside at t, in case of P_2. (4.31)

As long as the descriptions of the practices P_1 and P_2, that we have to supplement (4.31) with in order to make the picture complete, only resort to aspects that are readily available to a (normal) interpreter, this T-sentence may be thought to be capable of capturing the information that is required for interpretation.

However, various objections apply to this truth-theoretic semantic implementation. One problem that arises in this respect is that a preliminary identification of the practices in which an interpreter is involved, is not (always) available while adhering to the schema of radical interpretation. Rather, the practice in which an interpreter is immersed is identified in a piecemeal fashion, *during the interpretative effort itself*. This follows from the fact that we need to know (at least in broad outline) what a sentence means before we are in a position to grasp its purpose and function. Sentences evidently have a crucial purpose (or function) in linguistic practices; and intricate and complex practices are often (partly) linguistic. It is therefore clear that presupposing practices cannot be brought in line with the strictures of radical interpretation.

As far as the truth theory is concerned, what we call 'practices' is just a peculiar structure within the RHSs of T-sentences (i.e. conjunctions of conditionals). This may entice the hope that coming to gradually understand a practice will mirror gradual refinement of interpretative T-sentences. But observe that here, as in the rest of the process of radical interpretation, it is crucial that the recognition of a complex structure (e.g. a complex practice) must proceed along a path of simple steps. In Chapter 2 we saw why the incremental coming to meaning does work for interpretation: Atomic object language sentences are paired with relatively easy RHSs. Complex object language sentences are paired with complex RHSs *through* compositionally combining the relatively easy RHSs in a structural way. (Actually, the case for interpretation is every better than that: subsentential expressions are related to still easier parts.

For instance, the complex sentence (4.32), has a rather intricate RHSs (whether we can use the disquotational RHS is conditional on whether interpreter and interpretee speak the same language or not).

John walks and talks as if he knows a lot about foreign policy. (4.32)

But the intricate meta-linguistic description can be structurally reduced to satisfaction axioms, such as (4.33).

The reference of 'John' is John. (4.33)

f satisfies 'x walks' iff the referent of x walks.

But the same structural underpinning of complexity does not apply to the piecemeal generation of practice identification conditions. Let me explain why.

Another example sentence for which interpretation works just fine is: 'Sue saw the goose the mouse the cat the dog chased ate loved.' It is hard to match this object language sentence to the complex collection of environmental aspects that would, when described in the meta-language, provide the truth conditions that are needed in order to make the corresponding T-sentence come out right. Luckily, this sentence's subsentential components reappear in (a potential infinity of) other sentences. Moreover, many of these sentences will be of a less daunting complexity than the one presented above.[24]

What we are given in interpretation is (i) a number of utterances, and (ii) correlated characterizations of the situation in which these utterances have been made. As interpreters we then make the following two abstraction steps: (a) We extract (what I have called) recurrent uses from the disparate characterizations of the situations of utterance. (b) We look for a finite recursive specification that generates all these utterances. Crucially, these two processes must be performed so as to be in accordance with one another. It is because the utterances can be structurally generated, and because the identified recurrent uses are structurally related so as to mirror the recursive specification of the syntax, that we are able to construe a truth theory that specifies the right T-sentences (i.e. one in which (i) and (ii) are correctly related).

But in the case of meta-linguistic practice identifications we have nothing that comes close to this intricate interrelation of multiple (complex) identifications under a uniform axiomatic characterization. There is no recursive definition of practices to which practice identification descriptions can be related.[25]

The same problem can be formulated in a more casual vocabulary: it is not always possible to identify the practice at hand in terms of circumstantial evidence that is readily available to a (normal) interpreter. If I have no idea what practice I am immersed in (let alone a finite characterization that

[24] Observe that although in interpretative reality there are always simple sentences that contain the same subsentential expressions as complex sentences do, the availability of such simple sentences is not necessary for the argument. Due to the notion of meaning holism it is possible to arrive at an interpretative truth theory by just being given a multitude of complex sentences that contain shared subsentential expressions.

[25] The point may be somewhat difficult to appreciate and so an example might be helpful. In Section 2.3.2 we saw that, due to the notion of meaning holism, it was assured that non-interpretative but co-extensive meta-language sentences, like e.g. "'Snow is white' is T iff grass is green", could be filtered out because the subsentential parts had to be shared amongst a potential infinity of sentences. It was already suggested above that such a form of compositionally driven sub-part interrelation does not occur in the case of (descriptions of) practices. If my claim is valid, this would imply that if there are 'non-identical' but 'co-extensive' practices, then these will not be filtered out in a similar way.

And indeed, if I happen to perform one of the apple-sorting tasks discussed in Section 4.6.2 in a blue room, I might think that my partner's use of 'This apple is red' as conditional on the room's being blue. But it is evident that T-sentences that are supplemented with the blue-room-identification-condition for the practice at hand, will no longer work as soon as the room is being painted over in green. This does not prohibit me from instigating equally idiotic green-room-identification-conditions for a second practice. Observe how different this is from cases involving communicative contents!

I think uniquely lifts out that practice), then I just try to axiomatically specify a truth predicate that does the best job at interpreting the utterances I am presently confronted with. Under these conditions, interpretation may be perfectly possible even though I have no lights to go on as to what practice I am involved in.

A second objection to the present approach, ties in with the problem that is central to this essay, i.e. the problem of contextual dependency. If a single sentence can be used in a potentially infinite number of ways (appealing to an indefinite number of contextual aspects), and we give a single T-sentence for every object language sentence (with a meta-linguistic description that is supplemented in order to incorporate practices), then it is clear that we cannot explain context dependency in terms of universally quantified variables to which we relativize the truth predicate (since this would involve either an indefinite number of variables, or an indefinite number of values for these variables in response to which interpretation might change).

Therefore, if we want to explicate the contextual dependency of meaning to the extent that the intricate task of communication is elucidated, we cannot resort to the incorporation of individuation conditions for practices within the RHSs of T-sentences.

The problem could be circumvented by nothing that I am only confronted with a finite number of practices, and so will only have to relativize the truth predicate with respect to a finite number of contextual aspects. Although this is a valid solution (it is one big step in the direction of truth-theoretic contextualism, which discussed below), it is not in line with traditional conceptions of truth-theoretic semantics, in which the truth predicate is structured according to a uniform scheme (either 'is-T-in-*language*' or 'is-$T(S,t)$').

Luckily, there is an alternative way in which to frame the matter, one that I will call truth-theoretic contextualism. Instead of incorporating identification conditions for practices within the meta-linguistic descriptions, we want these identifications to be extrapolated from the process of truth theory construction itself. Under this alternative view we do not add extra things to the RHSs of T-sentences at all, but – instead– we bring about a more fundamental change: we do away with the idea that all the T-sentences an interpreter uses for interpretation must share a single truth predicate. Instead, what truth predicate she uses will depend on the practice in which she is involved.

These changes to the T-sentences that the theory generates result in the replacement of predicates such as 'is-$T(S,t)$' by 'is-$TP(S,t)$', where P names a practice. Instances of this for our apple-sorting example are (4.34) and (4.35).

'The apple is red' is-$TP_1(S,t)$ iff the apple is red on the outside at t. (4.34)

'The apple is red' is-$TP_2(S,t)$ iff the apple is red on the inside at t. (4.35)

Here the markers P_1 and P_2 are not indicative of the practices they represent, although we can make an additional specification, as e.g. in (4.36) and (4.37).

If my working partner cuts apples open, I use the truth predicate 'is-TP_1'. (4.36)

If my working partner scrutinizes the surface of apples, I use the truth predicate 'is-TP_2'. (4.37)

Observe that statements like (4.36) and (4.37) are just a theoretician's crutch. As far as the truth theory is concerned, it does not matter a bit how we would describe situations in which either one of the truth predicates is used. The only thing that is relevant from a theoretical perspective, is that in specific instances of communication, one of the predicates generates hypotheses that conform to instances of communicative interaction within the actual discourse situation, whereas the other predicate does not. That, from the viewpoint of the interpreter, it is possible to verbalize what practice is in vogue, provides an easy method for skimming through her set of available truth predicates. But this can only ease the process a bit, i.e. it is merely a heuristic (though potentially one that is quite useful in particular instances of linguistic interaction), it is not necessary in order for interpretation to proceed.

In interpretation the interpreter is confronted with conflicting truth conditions, and conflicts in truth conditions are resolved by instigating different truth predicates. How we characterize these truth predicates is something that may come in handy (again, as a heuristic) to the interpreter or the theoretician, in that it pares down the number of instances of communication that we need to consider in order to get at the right truth predicate that works for that discourse situation.

practices get instituted because of theory-internal considerations. But once a practice has been instigated, i.e. once an truth predicate has been added to the stock of truth predicates that an interpreter makes use of, there may be additional aspects of the context that help identify which of the existing truth predicates is to be deployed in a concrete instance of interpretation. (This is what we called a 'heuristic' in the above.)

An instance of such a heuristic, is that in the workspace (e.g. the office) I expect other recurrent uses to occur than when I'm at home. This may lead to my entering a conversation with the right truth predicate, i.e. the one that is tailored towards the right practice. So when picking up the phone I say either 'Office desk 75, how may I help you?' or 'Hello, Wouter speaking!', depending on readily available aspects of the environment (e.g. whether I am at the office or at home). The heuristic I use here may look like (4.38) and (4.39).

$$\text{If I am in the office, I use truth predicate 'is-}TP_{Office}\text{'.} \tag{4.38}$$

$$\text{If I am at home, I use truth predicate 'is-}TP_{Home}\text{'.} \tag{4.39}$$

Partitioning human linguistic interaction into practices means that different truth predicates work on different occasions. The individual truth predicates may be regarded as more or less well-tested modules for explaining packages of (partly) linguistic behaviour. As a corollary of this, truth-theoretic contextualism will generally slice uniform treatments of linguistic use much thinner than languages do.

As a side effect, the T-sentences that are now derived, because of their relative simplicity, provide much better hope for fitting the available data. The crucial step for the interpreter is to identify the right practices or truth predicates (identical concepts now!) on each and every use (and to look for a new practice, or a new , whenever interpretation breaks down). In the process we have moved from the task of finding *one very difficult fit*, involving a long and intricate RHS consisting of conjunctions of conditionals with nasty antecedents, to the task of finding *a multitude of very easy fits*, involving lean and simple RHSs.

Since different truth predicates characterize different practices, and since within a practice the contextual influence of meaning can be modelled in a finite way, the number of contextual aspects that we parameterize these truth predicates to need not always be the same. In relation to Kaplan [25]'s analysis, this means that we may have different contextual n-tuples for describing different practices.

In the light of the here elucidated theory, the strange examples that were given by Bezuidenhout, Recanati (and other contextualists and minimalists), were just so much stage setting in order to make us intuit different recurrent uses.

A description of an imaginative use makes us switch perspectives, it makes us see the actions of e.g. sorting apples or forecasting weather, in an altogether different light, i.e. under the auspices of a another truth-predicate than the one we would normally employ.

In my explication of truth-theoretic contextualism I do not distinguish between semantics and pragmatics (nor did Davidson in his formulation of). From this it follows that I am quite optimistic about the potential of formal approaches to elucidate those aspects of meaning than have traditionally been called 'pragmatic'. My view therefore conflicts with that of e.g. Cappelen and Lepore [10, p. 190], who formulate the following statement:

> There can be no systematic theory of speech act content.

According to my formulation of truth-theoretic contextualism, the RHSs of generated T-sentences will contain the same aspects of pragmatic enrichment as the LHS object language sentences do (provided speaker and interpreter speak somewhat the same language). The parity between the LHS and the RHS will always be established to a certain degree, but this just means that only those pragmatic effects that are relevant for communication to proceed/succeed will be included.

This does not mean that traditional pragmatic theories are useless or irrelevant. It only means that they do not contribute to an explanation of how communication is possible.

It is now clear that in my truth-theoretic contextualism the notion of a practice is not meant to obliterate the concept of truth at all (as opposed to some (mis)interpretations of Wittgenstein, say), but instead a practice *is* the concept of truth, as it is made manifest within communication.[26]

Since practices are now moulded into the various truth-predicates, the notion of a practice does itself not figure within a particular truth theory. It will therefore not be mentioned explicitly in any of the T-sentences the theory generates, and it will also have no empirical validation criteria that impinge on it in a direct way.

Neither is a practice to be thought of as an abstract measuring rod against which all our communicative endeavours will automatically start making sense. Practices are empirically validated in the same way (and against the same evidence) as truth-theories in truth-theoretic semantics are.

[26] A natural question that arises in this regard, is in how far the present use of the notion of a practice resembles Wittgenstein's treatment of the concept in his *Philosophical Investigations*. I take a practice to be a theoretical notion integrated into a truth-based theory of meaning, adapted in order to explain meaning's contextual dependence. This implies that the explication of meaning, relative to a practice, will have the form of a formal Tarski-style truth theory, whereas Wittgenstein uses the concept of a practice in order to dissolve the very notions of formality and theory (and maybe even truth).

The segregation into distinct practices is subject to empirical validation, since another way of slicing the truth predicate will fit the empirical data to a higher extent (avoiding over-fitting to occur[27]).

Since under the present approach the descriptions of practices are no part of the RHSs of T-sentences, it is clear that descriptions of practices do not themselves get formalized. This means that we do not have to introduce abstract terms, i.e. practices, within our meta-language. Instead, the present account tells us what the long though to be elusive notion of a practice concretely amounts to: it is emphtruth in use. Actually this description still contains some superfluous elements; we do not want to entice hope for something like *truth simpliciter*. The right way to phrase the matter is therefore, simply, *a practice is the concept of truth*, now who would have thought that?

[27]The notion of over-fitting is actually quite important in this respect, since it makes sure an interpreter ought not to cut the cake of truth in too many pieces that are too little to allow communication to succeed, see Section 5.5.

CHAPTER 5

COMMUNICATION

5.1 Contents

In Chapter 1 the principle of communication was introduced. According to this principle, a speaker can only communicate what an interpreter may be likely to catch up on. This means that the locus of communicative success does not exclusively lie in the speaker (or in the 'community' to which the speaker belongs). It can therefore be characterized negatively as that it seeks to abstain from the Humpty-Dumpty approach towards communication, i.e.: even though in instances of successful communication the meaning of an expression is identical to the intention the speaker had in uttering that expression, this surely does not imply that an expression means whatever the speaker takes it to mean.[1]

The set of contextual elements that speaker and hearer – once their theories have converged to a sufficient extent – share, can not be determined prior to the act of communication. An independent argument would have to be provided in order to show that only truth predicates that are parametrised with respect to a specific set of contextual aspects are fit for interpretation, whereas others are not. It is thus instances of communication that decide over the context-sensitivity of meaning.

The problem with the minimalism/contextualism-discussion is not that the notion of communication has not been discussed. As a matter of fact, in the monographs of Recanati and Cappelen and Lepore, communication plays an important role (and a decisive role in some of the key arguments there). But in the existing literature we see the pattern that, first, semantic meaning is distinguished from pragmatic meaning (this is often done by an appeal to intuitions, but never with an appeal to successful communication). Only after the dividing line between the semantic and the pragmatic has thus been established, do the arguments concerning communication set in. My approach has been rather different, since the goal of successful communication has been the driving force behind the discussion in the previous chapters (as well as in the present one). What is more, the division between the semantic and the pragmatic has not even been discussed for truth-theoretic contextualism as of yet. I do not think we need the division in order to get a clear view of the semantic

[1] For instance (5.1) can be uttered by the speaker while 'holding in her mind' the image of a woman, but this cannot reasonably be taken to be part of the meaning of (5.1) if the interpreter cannot be reasonably expected to know – or infer – that the speaker's account manager is a woman.

My account manager is a nice person. (5.1)

enterprise. The reason for this, is that if we perform the semantics task in the right way, we do not need pragmatics as a separate process. (A good semanticist prefers an empty garbage can.)

In this chapter the approach outlined in Chapter 4 will be compared to the alternative theories that both minimalists and contextualists have come up with. As a consequence of this, it will be clear that the present analysis – at least with respect to the topic of communication – is superior to rival approaches.

Although the various rival approaches all have their particular deficiencies, discussed in Section 5.2 the reason why truth-theoretic contextualism does suffice is that it circumvents the main problem that is present (either explicitly or implicitly) in all of the rival approaches. This main problem is discussed in Section 5.3.

In Section 5.4 we delve into the fundamental deficiencies that tie up with this main problem of existing theories of meaning. This dissection of the fundamental fallacies of existing accounts, will be given in terms of an interpretation of Davidson's distinction between prior and passing theories. The interpretation that I will give, will show it to be a viable alternative to existing accounts of communication.

In Section 5.5 the alternative approach towards communication, although circumventing the main problem, is shown to have a critical flaw when combined with Davidsonian truth-theoretic semantics. Luckily, my account of truth-theoretic contextualism, given in the previous chapter, fixes this flaw.

5.2 Problems with the various stances

In this section I briefly the problematic aspects of the various stances within the minimalism-/contextualism-debate, with respect to the elucidation of communication.

5.2.1 The problem with radical contextualism

According to radical contextualists the notion of compositionally is not central to the explication of communication (as it is according to minimalists and as it is in my account of truth-theoretic contextualism). The reason for this is that words are assigned 'semantic potentials' (an idea that derives from relevance theory). These semantic potentials provide the input for pragmatic processes that give as output the meanings of complex expressions.

Sometimes the semantic potential of words is thought to consist of nothing more than collations of applications of those words:

> The semantic potential of P for Tom at the end of his learning phase can thus be thought of as *a collection of legitimate situations of application*; that is, a collection of situations such that the members of the community agree that P applies in or to those situations.[2] [3] [45, p. 191]

But there are many problems with this approach. Most notably, the pragmatic processes that radical contextualists think should replace semantic analyses, hugely over-generate. There is not

[2]Recanati calls P both a predicate and a word, but in the use that he makes of the sign it can only be a word.

[3]The attempt to found communication in communal agreement, that is touched upon in Recanati's quote, will be discussed in Section 5.4.

even a respectable toy language for which pragmatic processes, operating on words along with their 'collections of legitimate situations of application', are able to generate the right meanings.[4] But I shall pose another objection to the radical contextualist attempt towards explaining communication. For in talking about (pragmatic) processes that operate on words that are associated with situations of use, radical contextualists have the matter reversed. According to the principle of meaning holism (introduced in Section 2.3.2, the assignment of meanings to words, in interpretation, can never precede the assignment of meanings to sentences.

Looking for relations between words and aspects of the world (e.g. 'legitimate situations of application') is a process that can only proceed provided sufficient headway has been made in uncovering the structure that is induced by a whole host of sentences (to which these words belong). Word-individuation proceeds by recognizing the role that parts of the phonological surface form play within the structural analysis of meaning. This implies that the identification of words can only be the *result* of a theory for interpretation.[5]

Communication evidently takes place on the sentential level.[6] And even in those sparse instances where a sentence consists of only a single word, the recognition that it in fact is one word that we are confronted with (and not e.g. two) can only be made antecedent to a primordial understanding of what that sentence means.

To give an example, 'gavagai' will be identified as two words as soon as 'kologai' gets paired to situations in which a running deer is signalled, and 'gavaro' gets paired to situations in which a rabbit is sitting still. It is only because no such situations occur in the examples given by Quine, that the sentence is taken to consist of a single word.

In truth-theoretic semantics, words are distinguished by the boundaries that the interpreter chooses to mark within the phonological surface structure, in order to make the thus identified components adhere to a compositional specification of meaning. The notion of a word is thus defined as the smallest entity that has a structural role to play in a truth theory that is used for interpretation.

Assuming the interpreter is in a position to draw these boundaries at the onset of the interpretative effort, is to violate the restrictions under which radical interpretation can proceed. Therefore, knowledge where the word boundaries are in a surface form, is something that gets established during the act of interpretation (and not prior to it).

Relevance theorists, and others who want to assign empirical content to pieces of language directly (i.e. without the need to posit an underlying compositional structure), are left with the problem that – on the one hand – meanings cannot be directly assigned to words, since it is a result of an interpretative endeavour what entities are to count as words; but – on the other hand – they cannot

[4]See [p. 225, 237]Stanley2005a for a critique in this line.

[5]Linguists of a structuralist bent may disagree in claiming that it is possible to identify word-boundaries by attending to syntactic properties only. Without wanting to digress in this delicate matter, I will merely remark that such structuralist accounts hinge on a primordial recognition of what is and what is not a syntactic string. But in language learning, syntactic intuitions are not learned prior to and independent of meaningful uses.

[6]Sometimes even sentences cannot be distinguished on the basis of the phonological surface form alone. In such cases, sentence individuation is just as conditional upon interpretation as word individuation is. (On another count, the distinction between word and morpheme-parts is, from a semantic point of view, often quite arbitrary. Take for instance 'played' and 'worked'. These may be translated as 'play some-time-ago' and 'work some-time-ago' in a language that distinguishes 'play', 'work', and 'ed' as individual words.)

assign empirical content to sentences (or multi-sentence chunks of discourse) either, since there is an infinite lot of them, so that making a finite number of meaning assignments will not suffice for general communicative purposes.

5.2.2 The problem with speaker intentions

According to Kaplan [25], even though the use of demonstratives is usually accompanied by objectively recognizable demonstrations, it are the directing intentions of the speaker that determine the referent of a token demonstrative. The objective aspects of a demonstration – such as pointings, glances, etc. – may help, or even 'determine', in a (quasi-)psychological sense, the interpreter's identification of the intended demonstratum. But the criterion of correctness for the interpreter's identification of the demonstratum is provided by the intentions of the speaker. So from a theoretical point of view, the objective aspects of demonstrations are semantically irrelevant (in the sense of non-essential) externalizations of internal states of affairs.[7]

But if we assume that the speaker is committed to the principle of communication, then he cannot be legitimately referring to an object that he knows (or may reasonably be taken to know) the interpreter has no access to. Since what a speaker intends is not generally available to the interpreter, a speaker who nevertheless thinks he can transmit meaning only conditional on his intentions (i.e. irrespective of the interpreter's being able to access the source), is like Humpty-Dumpty who thinks that words mean whatever he wants (or intends) them to mean.

This is not to say that speaker intentions can never be invoked in establishing meaning. It is quite obvious that speaker intentions are crucial to most if not all uses of indexicals and demonstratives. But the speaker has to make sure the hearer is in a position to distil his intentions.[8]

Once interpretation progresses, an interpreter comes to know a speaker by disseminating what the speaker takes his words to mean from what he believes. As a consequence, she is able to contextually resolve the meanings of his words with less need for explicit guidance. The speaker knows this, and as interpretation progresses he will adopt his behaviour accordingly.[9]

The present discussion allows us to contextualize the claim I made in Section 3.2.3, i.e. that there are only expression that appeal to subjective aspects of context (and none that appeal exclusively

[7] See also Borg [6, p. 162] and Recanati [43, p. 173-4], who hold similar views.

[8] In the early stages of interpretation, in which the interpreter knows (nigh to) nothing about the speaker, genuine uses of indexicality (and demonstratives) must be accompanied with explicit pointings (or other contextual aspects that can be identified by the interpreter). If certain identifications are not made by the interpreter (or are made in a wrong way), then the speaker ought to correct and further instruct the interpreter. If the interpreter got the clues regarding the resolution of context dependent expressions right, then this is evident from the way in which the communicative interchange proceeds. The speaker can observe a discongruence between the T-sentences the hearer employs and the ones that he himself makes use of, because the interpreter will make subsequent utterances under circumstances that the speaker may find inappropriate.

Digging one level deeper into the cycle of communicative interchange, a speaker's behaviour can be modelled as forming a theory about how he thinks the hearer is interpreting him. A speaker of this latter category (e.g. a teacher) is therefore not only able to reprimand the interpreter whenever she uses words in a different way, but is also in a position to make guided suggestions for steering her interpretative theory so as to be in closer correspondence with his own.

[9] Observe that the reason why the interpreter can now derive at the meaning of expressions with less explicit work on the part of the speaker, is – strictly speaking – not because the set of readily available evidence that she needs access to for interpretation to succeed has decreased. Rather, it is what she has access to that has been broadened so that the set of readily available evidence can be (partly) relocated (e.g. from explicit pointings to speaker intentions) through the incremental identification of what the speaker beliefs.

to objective aspects). I there used the personal pronoun 'I' as an example, and showed that – under certain circumstances – its reference could only be established relative to speaker intentions or subjective features. Now how does this remark relate to the present claim that in the early stages of (radical) interpretation an interpreter may not be able to recognize speaker intentions at all? For if at the same time subjective features are needed in order to uncover the meaning of *any* expression, then there is no way to enter the interpretative cycle.

But we must keep in mind that the positions that I was arguing against in Section 3.2.3 were all concerned with the segregation of expressions into either subjective or objective ones, *irrespective of use*. Now my claim was that for every expression there is always an (imaginary) use in which reference gets established relative to speaker intentions. But this does not mean that on *every* use the interpreter must have knowledge of the intentions of the speaker. More specifically, the linguistic uses that comprise the early stages of interpretation involve reference determinations relative to aspects of the readily available context that speaker and interpreter share. As interpretation progresses, reference determinations that exclusively rely on subjective features become a possibility.

From what I claimed in Section 3.2.3 and in the present section, we see that there is no such thing as a segregation of *expressions* in either subjective or objective ones; but there is such a segregation in *uses*.

5.2.3 The problem with perspectival contextualism

In Section 3.2.5 we discussed perspectival contextualism, according to which meaning can only be expressed by a sentence in its original context of use. Irrespective of the faulty reasoning that was already hinted at in that section, there is a more fundamental mistake that is central to this view.

If an utterance can only have meaning in context, and the meaning the sentence has in context cannot be (formally) specified irrespective of that context, then the transmission of meaning can only proceed if, together with the meaning, the context is transmitted too. But the central idea behind communication is that, at least in some cases, meaning can, to some extent, be transmitted across contexts of utterance. In other words, an account of human communicative abilities ought to at least explain how meaning can be transmitted without the need to adjoin the situated circumstance of utterance. This minimal requirement on theories of communication can never be met if the distinction between meaning and use never gets instigated in the first place, and yet this is what perspectival contextualists maintain.

5.2.4 The problem with minimalism

According to Cappelen and Lepore [10, p. 184-5] the minimal proposition is the only viable candidate for communicable content in language use.[10] The minimal proposition that a sentence expresses is claimed to be a (proper) part of the speech act content that is associated with any utterance of that sentence.[11] The particular role that the minimal proposition of a sentence plays in communication is derived from the fact that it is the only component that is identical to the speech act contents of

[10]The minimal proposition that is expressed by an utterance of a sentence is defined as the disquotation of that uttered sentence in which only expressions of the basic set have been contextually supplemented.

[11]The speech act content of an utterance is the set of propositions that are conveyed by the act of making that utterance. It thus includes both the minimal proposition that the semantic process comes up with, together with those propositions that the pragmatic process subsequently generates.

all the utterances of that sentence. This is what makes the concept of a minimal proposition *useful* for communicative purposes: it is that aspect of conveyed meaning that traverses all the contexts of utterance involving a single sentence.

In addition to being theoretically useful, the minimal proposition is also a notion that can be *realistically* derived within the process of interpretation. It can be easily identified by every competent language user because of the direct availability of the disquotation and because of the limited number of contextual supplementations that are required to be made (i.e. only for a handful of expressions that occur in the basic set).[12]

Moreover, speaker and interpreter know all these things. And because both speaker and interpreter know that the minimal proposition is what gets shared cross-contextually, the minimal proposition is that content which the speaker can expect the audience to grasp and expect the audience to expect the speaker to grasp. Also, it is that content which the audience can expect the speaker to grasp and can expect the speaker to expect the audience to grasp. It is also that content which someone who is not a participant of the context of utterance is able to grasp, e.g. when given a report on the matter. More specifically, it is the content which (both speaker and audience know) can be transmitted in instances of indirect quotation (or cross-contextual reproduction).

If the semantic content of a sentence is the set of propositions that are expressed by any possible utterance of that sentence, then it follows that the semantic content of every sentence is just the empty set. This is clear once we recognize that metaphor, irony, and deferred reference permeate language use, and can be introduced with respect to every sentence. If one wants to avoid such pitfalls, one has to show for every use of a sentence in which the purported minimal proposition does not obviously occur, that it is nevertheless present. Clearly, this is not easy to prove for the general case.

Borg [7, p. 352] notes that the minimal proposition will not generally be psychologically realized in both speaker and hearer, i.e. we have no reason to assume that speakers generally intend to assert such propositions, nor that hearers are generally able to grasp them. The same point is made by Recanati [42], who holds that the minimal proposition cannot be shown to have a psychological reality in the calculation of what is said by an act of utterance.[13]

But Cappelen and Lepore's position can be easily defended in the light of this critique by making a small alteration. We do this by noticing that the disquotation (supplemented for items belonging to the basic set) can be distilled from the uttered surface string irrespective of whether it belongs to the speech act content of the utterance involved. We thus do not need to define the minimal proposition in terms of what gets shared throughout the speech act contents of utterances of the same surface string.

Under the altered definition, the minimal proposition may still give some indication as to what the original utterance was (probably) about, even if it plays no role whatsoever in establishing what

[12]Provided each of these expressions appeals to a limited number of contextual aspects in a straightforward way; something that is not evidently the case at all, see Section 3.2.1.

[13]Do not be mislead by talk of 'psychology' here. Neither is it obvious that the minimal proposition is always *conceptually* necessary for the determination of speech act content.

is said (and, subsequently, in what is implicated). It is not a conceptual impossibility that what is said together with what is implicated gets computed independently of the minimal proposition.[14]

But if the computation of the minimal proposition is both psychologically and conceptually distinct from the computation of what is said, then it no longer plays a necessary role in the explication of communication. For the minimal proposition will often be something quite uninformative. Only very little of what gets communicated, if anything at all, must involve the minimal proposition (and thus only very little of communication is left to semantics).

Minimalists agree that minimal propositions need not express useful things (e.g. they may be trivially true or trivially false).[15] But Cappelen and Lepore [10, p. 165] claim that these minimal propositions are still useful in order to pare down the topic under consideration, thus providing the best possible access to the speaker's mind given the restricted knowledge that is available. For instance "John dances" is not about Peter, nor about walking.

But this is no longer clear under the above noted reformulation of the notion of a minimal proposition. For if the minimal proposition is no longer a part of the result of the generation process of speech act content, then there need be no connection at all between the content of the minimal proposition and speech act content. (E.g. it is not clear at all why the speech act content of an arbitrary utterance of "John dances" could not be about Peter's walking.) The 'best possible access to the speaker's mind'-defence of minimalism only works if the minimal proposition is conceptually related to what is communicated. But this is not the case.

From the sparse relevance (if any) that the minimal proposition is thus seen to have, [6] concludes that semantics has no role to play in the explication of communication. Cappelen and Lepore [10] think this conclusion is a bit too drastic, but they do concur that most of what gets conveyed in communication does not approbate formal treatment.

5.2.5 Similarity metrics

An alternative suggestion is that we can define successful communication in terms of the similarity between speaker's and interpreter's theories. In order to establish similarity between theories, we need a uniform metric on which to plot those theories. (This is irrespective of how we think a language user's competence ought to be theoretically represented.) Now this uniform metric must itself be something that is external to the interpretative meaning theories whose similarity is to be plotted (otherwise it makes no sense to appeal to a similarity metric).

From this it follows that the metric cannot make use of meanings (because then the whole point of introducing similarity vanishes). In radical interpretation we have such a metric in the measure of successful empirical confrontation (say, the percentage of verified T-sentence hypotheses). The measure is not absolute, since empirical confrontation can only comprise a marginal subset of the possible tests to which meaning theories might be put. Nevertheless, it does allow a rough approximation of similarity (an approximation that gets increasingly more accurate as the number of empirical confrontations increases).

[14] The minimal proposition can, for instance, play the role of fallback content even though it is not a (proper) part of any speech act content. This was already noted by Bach [2, p. 158] and Borg [7, p. 353].

[15] See e.g. Borg [6, p. 243]: "[...] the formal, modular semanticist is [...] in no way committed to claiming the liberal truth-conditions are what are communicated in these cases, nor even that the liberal truth-conditions are likely to be the subject of much interest to interlocutors in these cases."

Cappelen and Lepore [11]'s repudiation of the similarity metric is based on a very fundamental misunderstanding. They claim that only full identity between theories can and must be adhered to in communication. But remember that in their approach of semantic minimalism they presuppose that disquotational T-sentences (in which elements of the basic set have been contextually supplemented) characterize the interpretative endeavour. This means that object- and meta-language are assumed to be (exactly) the same.

That on Cappelen and Lepore's count it is possible to talk about complete coincidence in communication is therefore no surprise, once we observe that this position trivially follows from their assumption. But the assumption is itself a misguided one, since language-coincidence can obviously not be generically presupposed. Even if this were the case, the fact that two language users speak exactly the same language, i.e. that two users generate the same T-sentences for each and ever natural language sentence they encounter, is something that can only be found out *in communication*.

But since communication always consists of a finite number of utterances, the notion of complete coincidence cannot be made sense of within the framework of truth-theoretic semantics. As long as the verification procedure is *empirical*, as it indeed is, we can only hope to measure *likelyhoods* of coincidence.

But this does not mean that we can hope to measure the similarity between theories either. For in the case of similarity the same problems that were identified in the case of complete coincidence recur. The problem does not lie in the word 'complete' but in the word 'coincidence'. We cannot measure coincidence, whether complete or partial, in an absolute way.

5.3 The main problem

The main problem within the minimalism/contextualism-debate is one that is shared by both semantic minimalists and contextualists (either of a radical or a moderate bent). The problem surfaces e.g. in Taylor [53], when he asks:

> [...] just *why* we must specify *where* it is raining in the case of the present tense 'it is raining' if we are to express a complete proposition, but we need not specify *how much* it rained or for *how long* it rained in the case of the past tense 'rained' in order to express a complete proposition. [53, p. 90]

Taylor than answers this question in the following way:

> [T]he difference between ['It is raining'] and ['Laura danced the tango until she could dance no more'] depends entirely on the difference in the way in which 'to dance' and 'to rain' relate to the places where rainings and dancings happen. Unlike 'to rain', 'to dance' does not mark the place where a dance happens as the theme or undergoer of the dance. The theme or undergoer of the dance is the dancer herself. The place where a dancing "takes place" is merely the place where the dancer dances. When Laura is dancing in a place, it is not the place that undergoes the dancing. [53, p. 91]

Even though I do not understand the specifics of Taylor's reasoning here (I suppose in actuality these are *intuitions* he happens to have, that are presented as if they were *arguments*), what is important is the implicit assumption that the set of contextual parameters that is relevant for

uncovering the meaning of an expression is to be established *in a generic way*, i.e. independent of the use that we make of that expression, and independent of the peculiarities of the speaker and audience involved. According to Taylor, whether a parameter is needed in order to supplement an expression, is a characteristic *of that expression*; i.e. *an entirely linguistic affair*.

Even if the conceptual distinction that Taylor claims can be uncovered by appealing to the linguistic properties of an expression exclusively could be identified in a theoretically introspective way, then we are still confronted with the problem that the thus laid down analysis of meaning is entirely unrelated to instances of communication. If we want to explain how linguistic intercourse is possible then we need to relate abstract meanings to concrete situations, but this last step is never made.

Now I claim that the same critique applies to the work of all minimalists and all contextualists (that I know of). The key considerations that were distilled from the debate in Section 3.2, as well as the intuitions that were discussed in Section 3.3, are never related to instances of use in a systematic way. Both minimalists and contextualists (either implicitly or explicitly) think a semantic theory ought to be oriented towards a *language*. But in Section 4.7 we saw that in truth-theoretic contextualism the unit over which a semantic theory ought to range is a *practice*, i.e. a truth predicate that works for an interpreter to make sense of a speaker's words conditional on truth-theoretic considerations exclusively.

5.3.1 The minimalist response

There is a small number of participants within the debate that seem to escape the present critique. These are the semantic minimalists who maintain that semantics is not concerned with the explication of communication in the first place. All that I have said up till now depended on the assumption that semantics' role is the explication of communication. Therefore, it seems to be a legitimate move for minimalists to simply deny the validity of that assumption.

But there is a reason why I think no minimalist can discard the assumption that I make. If literal meaning (or semantic content) has been separated (i) from what is said, (ii) from the propositions that are communicated (i.e. what is implicated), and (iii) from instances of use, then we are left with no verification criteria whatsoever against which to check the validity of these purported literal meanings.

A minimalist semanticist thinks up truth-conditions that ought to make up the RHS of a T-sentence. How does he know that these are in fact the right truth-conditions? He cannot appeal to the communicative success that is brought about by a speaker uttering a sentence to a hearer. We cannot resort to a story about what the speaker intended to bring across in uttering that sentence. Neither can we make use of an account of what the hearer can reasonably be taken to have understood based on such an utterance. None of these things can play the role of a benchmark against which the correctness of a proposed semantic theory might be checked.[16]

Once communication is denied a constitutive role in semantics, mere intuition remains as a correctness-check on theory. This would not have been a problem had the intuitions regarding

[16]The same point is recognized by Elisabeth Camp when she says that the more we shift the burden of explaining what is said in an utterance from semantics to pragmatics, "the more indirect the empirical support for one's semantic theory becomes." She concludes that "[o]ur evidence for semantic theorizing must ultimately come from the utterances that ordinary speakers make and the interpretations that ordinary hearers assign to them." [8, p. 194].

the topic under consideration been universal; but this is clearly not the case (as is apparent from Section 3.3). A paraphrase of a well-known statement by Wittgenstein seems in place here: whatever the semanticist thinks is the literal meaning of a sentence *is* the literal meaning of that sentence. But this is not a position that can be legitimately termed scientific.

What we thus observe in semantic minimalism is what I would like to call an account of *detached semantics*. A semantic account is detached if there is no empirical validation of any kind with respect to which the correctness of theorizing can be evaluated. Since language itself is, first and foremost, an empirical phenomenon, this is a particularly unfortunate property for a theory of meaning to have.[17]

5.3.2 No use

Minimalists seem to be perfectly at ease with the fact that their theories cannot be empirically confirmed. E.g. Borg sees the semanticist's business as perfectly isolated from instances of language use:

> To capture the literal meaning of a natural language sentence it is necessary and sufficient to determine its truth-conditional or propositional content, and, the formal theorist claims, this content can be recovered via syntactic trails alone.

Now the first claim is right, and is adhered to in my formulation of truth-theoretic contextualism: the meaning of a natural language sentence *is* that sentence's truth-conditional content. However, the notion of truth, as it occurs in natural language analysis, itself depends on use (as was explained in Section 4.6 and Section 4.7).

But with the second claim I cannot agree. Semantics is all about those aspects of meaning that are formalizable. What is formalizable are truth-conditions, and these are defined by the satisfaction relation.[18] But this does not imply that we ought to focus on the 'repeatable' and 'code-like features' exclusively.

[17] Cappelen and Lepore [10] seem to recognize this. They therefore introduce what they call 'the mistaken assumption':

> A theory of semantic content is adequate just in case it accounts for all or most of the intuitions speakers have about speech act content, i.e., intuitions about what speakers say, assert, claim, and state by uttering sentences. [9, p. 278]

Now they suggest that their tests for context-sensitivity circumvent this mistaken assumption by focussing on the speech act content that a wide range of utterances of a certain sentence have in common (see [10, p. 57]). From this perspective it is especially unfortunate that the tests for context-sensitivity yield such ambivalent results (see Section 3.3.1 and Section 3.3.2). For if there is no consensus regarding the speech act content of an utterance, then neither is it possible to determine which aspects of content are shared (yet this is what their purported solution to the mistaken assumption requires).

The problem is solved as soon as we see the matter within the broader framework of truth-theoretic semantics, where the requirement that only aspects that occur in a wide enough range of utterances is replaced with (or implemented by) the requirement of meaning holism.

[18] All the other, i.e. non-truth-conditional aspects of meaning, ought to be excluded from the realm of semantics, they go by such inexact names as 'Färbung' and 'Beleuchtung', or 'Kraft'. (It appears that the analysis of these other aspects of meaning has not advanced beyond the sparse treatment they originally received in Frege's work.)

There is no way in which any meaning theory can be 'code-like'. A linguistic sign can only have meaning *for someone* and *within use*. In terms of string-manipulation, however intricate the task performed may be, we can only make sense of syntactic derivations.[19]

The reason why a purely syntactic approach to meaning is not available when it comes to natural languages, is that even though the theory itself may be defined by formal means, the processes of T-sentence construction and the empirical validation of existing T-sentences, both rely on the human ability to interrelate the formal linguistic structure to aspects of the world, under a specific (and more or less stable) use.

5.3.3 Summarizing

So much must be clear by now: if we remark that "The cat lies on the mat" (see Section 3.3.3) has its original meaning only when we find ourselves near the surface of the earth, then we find ourselves in the business of specifying truth-conditions for any conceivable context. But this is not what language is meant to do; it ought not to be viewed as a tool that extends use indefinitely. The meaning of the sentence is perfectly clear as long as it is used within the sphere of praxis under which it was instigated.

Both speaker and interpreter, when communicating, find themselves within the same sphere of praxis, and their words cannot reach outside of that sphere while still retaining their original meaning. Of course it is always possible to widen the sphere and to extend the praxis. But fixing the meaning of an expression across all (or many) uses, is itself a very specific property that is founded on a very specific use. That cross-contextual disquotation is often possible, is due to the human capacity to imagine a possible situation and a possible practice in which the disquoted sentence might have been uttered.

5.4 Communication and convention

In the previous section I argued that a language is not the kind of thing a semantic theory ought to be oriented towards. This was already implemented in my truth-theoretic contextualism, in which practices took over the role of characterizing truth predicates. Moreover, characterizing truth predicates with practices was shown to solve the problem of the contextual dependency of meaning. But the inability of languages to characterize truth-predicates in the way that is relevant for interpretation has not been sufficiently refuted.

Therefore, in Section 5.4.1 I explain why a language is not the right concept for framing the data set a semantic theory ought to explain. In Section 5.4 I then proceed in more general terms, and show that from the point of view of communication, the notion of a language is just as useless as it is for characterizing interpretative truth predicates. The idea underlying the communicative purpose of the notion of a language, is that communication can be explained in terms of a shared linguistic competence that the speech participants possess. Learning a new language consists of gaining this same shared linguistic competence. Finally, in Section 5.4.3 I will describe an alternative to such

[19] In formal languages, e.g. those for which Tarski provided truth-based theories of meaning, string manipulation suffices because of the syntactic nature of the semantics of formal languages. When we shift to natural languages, such a syntactic approach becomes insufficient.

conventional or competence-based explanations of communication. The alternative position is a specific interpretation of Davidson [21]'s distinction between prior and passing theories.[20]

5.4.1 The notion of a language under scrutiny

Let us take 'English' as an example. As far as it designates a language, it designates an infinite number of sentences that can be generated from a finite base (this is the standard definition in linguistics). But the description 'the English language' in actuality designates an indefinite number of such theories (or axiomatic bases), each with a potentially infinite generative capacity. This means that I can start using the sentence "The apple is red" in order to convey that at least half of the intended apple's surface is red, or that it is red on the inside, etc. So I can convey a multitude of different things and still be taken to speak (some form of) English.

The traditional goal of semantics is to give the meaning for all linguistic behaviour that may be termed 'English'. In Section 4.7 we saw that the truth predicate we use for interpretation is not in sync with what we call a natural language. This means that no natural language can ever be identified in terms of a *single* axiomatic base.[21]

This is only evident once we recognize that natural languages are generally identified for extra-linguistic reasons, e.g. based on ethnic, social, economic, or geographic distinctions. Whatever the identification conditions for specific languages may be, it is clear that 'a language' is not a notion that was invented by semantic theoreticians. It would therefore not be surprising if the term at some point turns out to be no longer suitable for the purpose of semantic theorizing.[22]

In Section 4.5 we saw that even for a static set of sentences, the meanings that are assigned to those sentences change through changes in use. Since the meanings for the bereaved sentences derive from the truth theory that generates the right results for the sentences that are in active use, it is clear that if the set of actively used sentences changes over time (or changes intra-personally), the meanings of the bereaved sentences will be altered accordingly.

Tarski [52] proved that it is there can be no monolythic concept of truth (that can be formulated in a consistent theory). Combine this with the fact that a natural language creatively allows its users to express a multitude of things, and it follows that truth can never be defined for a natural language.[23]

Even though [52] showed that it is not possible to define one truth predicate for all formal languages, this did not invalidate the process of defining a truth predicate, as long as the predicate was geared towards a single formal language. We saw that natural languages are so open-ended that they can potentially encompass every formal language. If we combine this observation with Tarski's proof, under the Davidsonian transformation described in Section 2.5, we understand that there cannot be a monolythic truth predicate that works for the interpretation of an entire language.

[20] The interpretation I give is not a piece of exegetical work, and conflicting interpretations of the same distinction exist, most notably in [28].

[21] Try to define a uniform truth predicate for the two uses of the sentence "The apple is red" in Section 4.6. You can always come up with yet another use of that same sentence, that is orthogonal to the predications your truth predicate currently gives.

[22] It is characteristic of the maturity of a scientific discipline that concepts that derive from everyday language use get replaced by new terms that are more specifically tailored towards the investigative purpose.

[23] In other words, a singe natural language's creative potential encompasses all formal languages.

Parallel to Tarski's undertaking, this does not invalidate the use of truth predicates for explaining interpretation. But it does mean that the everyday concept of 'a language' is not specific enough to do the work of partitioning the universal concept of truth into interpretative truth predicates in the way that is required for successful communication.

5.4.2 Why convention does not suffice for communication

The idea might be that although traditional characterizations of natural languages may not suffice for theoretical purposes, as was illustrated in Section 5.4.1, the explication of meaning still requires a shared linguistic competence. This idea is prominent in the work of Cappelen and Ernie LeporeLepore. This is why they think disquotation suffices for giving valid T-sentences: because interpreter and interpretee are assumed to share linguistic competence.

It is a truism that speakers do as a matter of fact deviate – either intentionally or inadvertently – from communal norms, while their hearers are still in a position to distil what was intended to be conveyed. But my present claim is much stronger, namely that knowledge of linguistic conventions cannot ever suffice for interpretation. The idea is that conventions underdetermine potential instances of situated response.

We can, rather theoretically, imagine a situation in which two people hold absolutely identical meaning theories.[24] When I say that two truth theories are identical, I mean that the totality of (potentially) generated T-sentences is identical, because it does not make sense to talk about identical axiomatic bases. The reason why talk f the latter does not make sense, is that there are multiple specifications of the satisfaction predicate that all generate identical T-sentence sets. Now even under this supposition, i.e. that two people hold generatively identical truth theories, we must still remember Davidson's admonition that "the problem of radical interpretation is domestic as well as foreign: it surfaces for speakers of the same language in the form of the question, how can it be determined that the language is the same?"[25]

If two people in fact have identical meaning theories, then this fact need not be epistemologically accessible to either one of them. In practice, the identification conditions for languages (and, thus, the verdicts regarding language similitude) are satisfied after T-sentence similitude has been establishing on a finite number of concrete instances. (If we take another dirty aspect of concrete linguistic interaction to heart, then we might even formulate language identity in terms of relatively often coinciding T-sentences.) It is an essential characteristic of linguistic practice that the identification conditions for a language, even when thus loosely defined, cannot themselves be conventionally established.

The fundamental reason as to why convention falls short of explaining communication, is that in the process of interpretation belief and meaning are taken to be fundamentally interlocking concepts (see Section 2.4.1). If belief is counted as the slack that exists between what is true and what is thought (or claimed, or maintained) to be true, then it is evident that *even utter agreement in concrete linguistic interaction cannot imply meaning theory identity.* There will be many alternative theories in which both meaning and belief are taken to vary. So that even if meaning were to be

[24] One can postulate a whole community of speakers employing the same meaning theory here, if one feels that it strengthens the argument.

[25] From [17, p. 125].

determined according to convention, empirical evidence could never corroborate this. In other words, evidence applies to both meaning *and belief*, where the distinction between the two components cannot itself be conventionally established.

This is the reason why knowledge of conventions cannot be a necessary ingredient of a theory of meaning.[26]

5.4.3 Communication and what gets shared

We saw in the foregoing that communication was defined (loosely) as the interpretative endeavour in which speaker and hearer, though they might start out at different positions, eventually come to coincide with respect to the meaning theory they employ, with respect to the communicative exchange that is under consideration.[27] In this section I will present an interpretation of the diction that is made in Davidson [21] between prior and passing theories.

In communication, the starting positions – of both speaker and interpreter – are characterized by their *prior theories*. For an interpreter, the prior theory will characterize her disposition to interpret the speaker prior to the onset of the communicative exchange. For a speaker, the prior theory will characterize what he believes the interpreter's prior theory to be.[28]

At the point at which communication becomes successful (becomes genuine communication!), speaker and interpreter are said to employ a *passing theory*. For the interpreter this characterizes her disposition to interpret the speaker's utterances in the midst of communication. For the speaker it characterizes the theory he intends the interpreter to use. It is thus the passing theory that expresses what the speaker's words mean. And it is only when speaker and interpreter employ the same passing theory that communication is said to be successful.

Both prior and passing theories have the character of a Tarski-style truth theory turned into a meaning theory (as described in Chapter 2). What is different now is that these theories are taken to model the dispositions of a language user with respect to a single speaker, and with respect to a bounded time interval (encompassing the communicative exchange). Because the characterization of the truth predicate consists of a speaker and a time interval, a meaning theory is seen to be a transient construct that changes in the intermittent flux of interpretation.

Framing communication in this way, interpretative success comes to depend solely on a speaker's and an interpreter's passing theories' coinciding. Since this condition for success does not state anything about prior theory similitude, what prior theories speaker and interpreter are using is theoretically irrelevant to the explication of communication (though it may be a useful crutch in practice).

[26] The analysis in this section opposes that of Lepore and Ludwig [28, p. 270-1].

[27] This last qualification is needed, since for the success of any communicative act it may not be relevant that speaker and interpreter agree on the meanings of all the expressions in their language. Rather, it seems more likely that interpretation regarding the sentence "Snow is white" can already be successful when both speaker and hearer share the same meaning theory for sentences containing the words 'snow', 'is', and 'white'.

[28] I agree with Lepore and Ludwig [28, p. 273] that this is a *first order theory* that characterizes the speaker's disposition to use his words conditional on his wanting to use those words in accordance with how he would suppose the interpreter is interpreting them on the onset of the act of communication. This interpretation emphasizes that neither speaker not interpreter must be claimed to have anything like explicit knowledge of the theories philosophers are keen to attribute to them. More specifically, the interpreter is not required to have active beliefs regarding the interpreter's prior theory (i.e. a second-order theory).

The prior theory is taken to model knowledge that would suffice for an interpreter to understand the meaning of an indefinite number of expressions. Linguistic competence, if it exists, ought to be located within prior theories. Combined with the earlier point that prior theories need not coincide, we conclude that in the theory of communication that Davidson proposes, there is no such thing as a shared linguistic competence. In other words, there is not something that both speaker and interpreter know (in the loose sense of 'know') and that is necessary (and/or sufficient) for linguistic communication.[29]

What this alternative theory of communication amounts to, is that the act of interpretation is no longer taken to be conditional on the speaker's conforming to a communal standard. I think some interpreters of Davidson have not quite coincided on his intended passing theory regarding this latter point. According to Lepore and Ludwig, Davidson's denial of there being anything like a language, still "leaves plenty of room for thinking that the prior and passing theories, for most interpreters and speakers in a linguistic community share a lot in common, certainly enough to make sense of the idea that they share a language."[30] Now this is true in some sense of course, but what I take to be the locus of Davidson's argument is that although speakers oftentimes employ meaning theories that are compatible, in the sense that interpretation unproblematically proceeds, this does not imply that there need be anything that gets concretely shared in the process. What goes wrong in correlating communicative success in linguistic practices with the sharing of content between speakers, is indicated by what it is that Lepore and Ludwig think gets shared:

> We can think of this [= speakers of a linguistic community sharing a lot in common] simply as that shared subset of axioms that characterize their prior and passing theories: given any overlap, they share a language. [28, p. 276]

This passage is fundamentally flawed. One cannot think simply of sharing meaning theories in terms of sharing sets of axioms. As was already noted above, for any one meaning theory, many axiomatic definitions of the satisfaction predicate would equally suffice. What gets shared, if anything gets shared, is T-sentences and relations between those T-sentences. But since there is an infinite lot of them, we are stuck with the problem that *what can be compared is infinite, while what is finite cannot be sensibly compared.*

There is thus no way in which meaning theory coincidence is to be measured, unless through the interpretative effort itself. But this implies that the 'location' of linguistic meaning has already been shifted. Since if what gets shared between two speakers can only be borne out in the particularity of the communicative act, then a communal conception of meaning does simply not apply. Formulated in slightly different terms, we must conclude that there is no sensible attribution of linguistic competence that might aid in an act of communication prior to the occurrence of that very act of communication.

That another person and I share a language / are to be attributed similar linguistic competences / can be counted as adhering to communal standards, can only be found out in communication.

[29] Davidson originally formulated his finding as that "there is no such thing as a language", at the end of Davidson [21]. The formulation that I employed above is a bit more accurate, i.e. there is no such thing as shared linguistic competence existing between people that are said to speak the same natural language.

[30] Lepore and Ludwig [28, p. 276].

That is the ultimate reason why radical communication is native as well as foreign, i.e. since being part of a linguistic community has no role to play in an exposition of communication.

Another effort of Lepore and Ludwig to localize some notion of linguistic competence within Davidson's distinction between a prior and a passing theory, is characterized by the following passage:

> [...] the theory that characterizes our dispositions to interpret someone as, so far as we can tell, an ideal speaker of the public language. The theory characterizes our competence in the public language. To the extent that it corresponds to the public norms, we can be said to be competent in the public language. Lepore and Ludwig [28, p. 274]

The appeal to an ideal speaker might be thought problematic, but it can be defined (rather loosely) as a speaker who is generally thought of as belonging to a linguistic community, and who has dispositions that are conditional on what he believes the words in the linguistic community mean. Under certain conditions (that can be summarized in the characterization that the speaker is competent in the language under consideration) we can reasonably assume that the dispositions of such a speaker can to a vast extent be characterized by a limited set of abstractions, in the sense of generalizations, over (a subset of) the members of that linguistic community.

Let us take a *generalized theory* to be the theory an ideal (or normal) speaker of the language would speak. This generalized theory is then taken to represent the language according to public norms.

But there is no reason why the generalized theory – though it embodies the prior theory that a speaker may be likely to use in case of encountering another speech participant for the first time – has to be the same for different people. Without additional conventions as to how one is to construct a generalized theory out of the many specific meaning theories one has come across in concrete instances of communication, the idea of cross-personally sharing a generalized theory cannot be assured. The case is particularly problematic, since the collection of meaning theories that one has come across in the past is never (exactly) the same for two different people.

The only way in which the concept of a generalized meaning theory can do its work, is by conventionally determining how one is to construct such a generalized meaning theory. It therefore looks like, on Lepore and Ludwig's proposal, preliminary convention breeds linguistic competence, but this is of course highly circular.

I conclude that even though it is in principle possible to distil for every language user – from his or her history of past interactions with others – a generalized theory, this does not make that theory a shared one. (Generalizing and sharing are not the same.)

5.5 Semantic solipsism as a threat to truth-theoretic contextualism

In Section 5.3 we saw that the main problem in the minimalism/contextualism-debate has been an excessive reliance upon the obscure notion of a language. In Section 5.4 I explained that a semantic theory that revolves around the notion of a language will not be able to explain communication. In Section 5.4.3 my interpretation of Davidson's distinction between prior and passing theories was presented as an alternative to competence-based explanations of communication.

This alternative theory of communication is compatible with semantic theories that do not make use of the notion of a language. Instances of such compatible theories are Davidson's truth-theoretic

semantics and my reformulation of that, as it occurs in Section 4.6 and Section 4.7, i.e. truth-theoretic contextualism. However, when we combine the alternative theory of communication with Davidson's truth-theoretic semantics, we are confronted with a fundamental problem.

The problem with truth-theoretic semantics, once we characterize the truth predicate with speaker/interval pairs (as in Section 2.5.1), is that the empirical data that is available to the interpreter, and that answers the strictures of the speaker/interval pair's values, may not suffice in order to verify the truth theory. If we can only consider a single speaker, and if the time interval is too short, we run out of empirical evidence with which to verify our theory.

I will call this the risk of *semantic solipsism*, because the danger arises that the theory the interpreter has come up with cannot be sufficiently geared towards the speaker's speech, in order to make communication possible. Under the sparsity of empirical instances of linguistic interaction, language users would be effectively locked up within their private meta-language (i.e. one that is never structurally related to an object language to the extent that is required for communication to proceed).[31]

The problem of semantic solipsism seems insurmountable in the work of Davidson, because it is unclear in what way previously formulated truth theories can be reused for interpreting new linguistic interactions. Because of the way in which the truth predicate is characterized, i.e. by a speaker/interval-pair, it is fixed to a specific point in both space and time. This means that for every interpretative encounter the interpreter will have to reset her interpretative undertaking. In other words: all interpretation will have to be radical interpretation. And what is more: all interpretation will have to be radical interpretation with an extremely limited number of verification instances, on the basis of which the right truth predicate must be found.

Of course the problem can be solved by assuming that some parts of older meaning theories get carried over into the new interpretative endeavour. But as soon as it becomes *necessary* to import (parts of) previously invented theories of meaning for present meaning theory construction, this only shows that the speaker/interval-pair is not the right thing to characterize truth predicates with.

As a matter of fact, it has always been unclear why Davidson thought the truth predicate should be geared towards a single speaker. There are ample instances of linguistic communication that involve multiple speakers, together shaping the interpreter's truth theory. It was also unclear what time interval ought to be taken as the locus of the truth predicate. A fixed number of seconds perhaps?[32]

And yet Davidson did not choose his parameters arbitrarily. It is not entirely coincidental that practices are often identified in both temporal and locational terms, and/or by identifying the participants of the practice, most notably comprising the (main) speaker. The only problem is that

[31] One may think of another problem in this respect: the individual truth theories (for all the individual speaker-/interval-pairs) that an interpreter might come up with need not have anything in common. I will call this the property of *semantic eccelecticism*. If an interpreter switches truth predicates, she will possibly be interpreting utterances in a radically different way. I do not think that semantic ecclecticism is is a problem for a theory of meaning though. In most cases the sketched situation will not arise, since the practices in which the interpreter gets involved will not be all that different. Constancy of behaviour is induced by constancy of environment.

[32] Observe that there are problems with identifying the time interval with the duration of the communicative interchange. We would have to keep the truth predicate steady even under quite erratic switches in behaviour on the speaker's side. Instead, what we want is a criterion of time interval length that itself depends on the recognition of such erratic behaviour.

Davidson has not given any theory-internal consideration as to why and how these two parameters must be used.

But in truth-theoretic contextualism we do find theory-internal considerations that lift out a practice as the correct characterization of the truth predicate. Because practices may involve multiple speakers and interpreters, because the time interval is now indicated, and because a practice never occurs only once, but recurs (and most of the time recurs very often), the problem of semantic solipsism does not arise (or is at least much less likely to arise) in truth-theoretic contextualism.

The set of practices a person is involved in some of the time, can best be though of as a concise library of truth predicates. Each of these predicates has certain preliminary identification conditions attached to them. The identification conditions that are attached to a truth predicate are never necessary nor stringent. They merely function as a heuristic in finding the optimal truth predicate for describing the current linguistic interaction (this was already explained in Section 4.7).

A speaker first picks out the truth predicate that most optimally fits the data. This may involve as little data as only one utterance (as in the apple sorting case discussed in Section 4.6). For more intricate cases, a larger number of utterances may be needed in order to lift out a unique truth predicate from amongst the alternatives that are still in line with present observations.

In interpretation it is possible, and advisable, to consider the measure of fit of a number of truth predicates in parallel. In line with this, the computational burden that is put on the interpreter is highest at the onset of an instance of linguistic interaction (all truth predicates that adhere to the heuristic description of the situation are being considered in parallel), and is lowered after each subsequently observed utterance. If the right truth predicate was not included under the initial heuristic description, the interpreter will have to return to her original library of available truth predicates (with either a different or no heuristic to go on).

If a speaker is unable to determine a heuristic description of the situation – or she discovers her initial characterization of the heuristic, in combination with observed instances of utterance, made her end up with no truth predicate whatsoever (i.e. the empty set) –, the processing power that is required in order to make her find the right truth predicate that does suffice for communication will increase.

If a speaker ends up with no truth predicate that stands the empirical test against the instances of the current linguistic interaction, she will take the one that applies most nearly (i.e. the truth predicate most of whose T-sentence hypotheses get confirmed), and she will try to augment that existing truth predicate's precision by altering its axiomatic structure so as to be in accordance with the current interaction.

(Observe that the computational burden under these various circumstances of linguistic interaction seems to be in line with common sense intuition.)

CHAPTER 6

CONCLUSION

In this essay I hope to have shown how the minimalism/contextualism dispute can be resolved. There is actually not a fundamental opposition between use-based theories of meaning and the more traditional, formal approach towards semantics.

A similar opposition (if not the same one) can be observed in philosophy more broadly. It is the opposition between those who take the notion of truth to be central to philosophical theorizing, and those who think the notion of truth is not (or maybe no longer) relevant. Philosophers of the latter denomination often seek to replace the traditional notion of truth with the notion of 'what works'. Since what works depends on the situation in which we find ourselves, as well as on the intentions of the people who are involved in that situation, I sense more that a similarity between the concept of semantic contextualism and pragmatism.

This opposition between traditionalists and pragmatists ought to be reconciled, just as the opposition between minimalists and contextualists has been reconciled in the present essay. The opposing sides in this dichotomy were shown to be both problematic in their own respect. The problem with the traditional account of truth-theoretic semantics is that it relies on principles that are not underwitten by the truth theory itself (this concerns both the notion of a language and the characterization of the truth predicate by a speaker/interval-pair). The problem with the contextualist account is that it suggest the notion of 'instances of use' and of 'what works' do not need to be structured in anything like the formal sense. But human (linguistic) behaviour is structured to a daunting extent. There is no inherent argumentation available why the notion of *structure* must dissolve once we move from monolithic truth to *use* or 'what works' (or once we move from the minimal proposition to speech act content).

In devising truth-theoretic contextualism as the right model for semantic theorizing, I tried to unify Tarski's refutation of the viability of a monolithic truth predicate, as well as his formalization and material adequacy criteria, with the sensible bits that can be drawn from the contextualist's critique on the abstract truth-theoretic approach towards meaning (abstract in the sense of not geared towards truth). I also hope to have shown, in my solution to the problem, that what is misguided about these latter theories is not that they revolve around the notion of truth, but that they are abstract. In other words, minimalists *do not put truth to use*, whereas truth is perfectly suited to do just that!

There are still some rough edges about the present undertaking that will need to be addressed in future research. A major topics has already been tapped on above. It is the question how the here noted identity between practices and truth predicates should be placed within the broader discussion

involving pragmatism. This especially holds true for the work of Rorty I guess, but also the way in which the present results interact with the work of William James and John Dewey may be thought interesting. As I have explained, I see nothing radical in 'replacing' the notion of truth with the notion of 'what works'. We are trading triangles for three-sided polygons here.

And so I predict that as soon as we want the notion of 'what works' to play an elucidating role in explaining the complexity of human (linguistic) behaviour, it will be seen to coincide with the notion truth. The other way round: once the notion of truth, as it figures in the elucidation of complex systems, gets truly geared towards instances of situated use – something the contextualists have made apparent that should be done for the investigation of the complex system of language interchange – it will be seen to coincide with the notion of 'what works'.

Another problem that I have not touched upon in this essay concerns the relocation of linguistic normativity. Some work in this area is needed in order to underwrite the changes to the semantic endeavour that have been proposed here. One specific aspect that I would have elaborated upon if considerations of space would not have constrained me, is the relation between the way in which I defined a practice in this essay, and the way in which Wittgenstein makes use of the notion in the *Philosophical Investigations*. This aspect is interesting, since in Section 5.4 I shunned community's role in the determination of meaning. But on Wittgenstein's count the community of language users seems all-important in establish linguistic normativity (or at least this is the case for Kripke's interpretation of Wittgenstein).

Besides these suggestions for future research, I have just as many (if not more) suggestions for research that should *not* be conducted in the future. So no semanticist nor pragmatist should base his or her theory upon a notion of propositionality (or a lack thereof), without explaining what a proposition is in the first place. Also, no semanticist should publish research that is based upon intuitions that are inconclusive and that do not contribute to a theoretical insight. Also, researchers must not loose sight of the crucial distinction between logical structure and conceptual analysis. Not every aspect of conceptual analysis is relevant for communicative purposes. Therefore, not every shade of a conceptual distinction should be left to semantics. (There are many shades of conceptual distinction that belong to e.g. epistemology or ethics.) Finally, no minimalist should be allowed to presuppose language identity between speaker and hearer, without further argument. It is clear that the disquotational approach towards meaning is illegitimate. The only thing that legitimates a disquotational account of meaning is a speaker and hearer's sharing the same truth theory. But truth theory identity is something that can only be established as the result of successful communication (making the account perfectly circular).

BIBLIOGRAPHY

[1] J. Austin. *How to do Things with Words*. Oxford University Press, 1962. <19>

[2] K. Bach. Conversational implicature. *Mind and Language*, 9:124–162, 1994. <35, 73>

[3] K. Bach. Semantic slack: What is said and more. In S. L. Tsohatzidis, editor, *Foundations of Speech Act Theory*, pages 267–291. Routledge, 1994. <35>

[4] K. Bach. The semantics/pragmatics distinction: What it is and why it matters. In K. Turner, editor, *The Semantics-Pragmatics Interface from Different Points of View*, pages 65–84. Elsevier, 1999. <25>

[5] A. Bezuidenhout. Truth-conditional pragmatics. *Philosophical Perspectives*, 16:105–134, 2002. <33, 34>

[6] E. Borg. *Minimal Semantics*. Oxford University Press, 2004. <1, 23, 24, 25, 29, 32, 36, 48, 70, 73>

[7] E. Borg. Minimalism versus contextualism in semantics. In G. Preyer and G. Peter, editors, *Context-Sensitivity and Semantic Minimalism*, chapter 14, pages 339–360. Oxford University Press, 2007. <22, 23, 33, 72, 73>

[8] E. Camp. Prudent semantics meets wanton speech act pluralism. In G. Preyer and G. Peter, editors, *Context-Sensitivity and Semantic Minimalism*, chapter 7, pages 194–215. Oxford University Press, 2007. <75>

[9] H. Cappelen and E. Lepore. On an alleged connection between indirect quotation and semantic theory. *Mind and Language*, 12:278–296, 1997. <76>

[10] H. Cappelen and E. Lepore. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Blackwell Publishing, 2005. <22, 25, 28, 33, 34, 35, 36, 43, 48, 55, 60, 65, 71, 73, 76>

[11] H. Cappelen and E. Lepore. Relevance theory and shared content. In N. Burton-Roberts, editor, *Advances in Pragmatics*. Palgrave Macmillian, 2007. <74>

[12] R. Carston. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Blackwell, 2002. <33>

[13] E. Corazza and J. Dokic. Sense and insensibility or where minimalism meets contextualism. In *Context-Sensitivity and Semantic Minimalism*, chapter 6, pages 169–193. Oxford University Press, 2007. <28>

[14] D. Davidson. Truth and meaning. In *Inquiries into Truth and Interpretation*, chapter 2, pages 17–42. Clarendon Press, second edition, 2001. <7, 9, 11, 16, 18>

[15] D. Davidson. On saying that. In *Inquiries into Truth and Interpretation*, chapter 7, pages 93–108. Clarendon Press, 2001. <9, 10, 47>

[16] D. Davidson. Semantics for natural languages. In *Inquiries into Truth and Interpretation*, chapter 4, pages 55–64. Clarendon Press, 2001. <18>

[17] D. Davidson. In defence of convention t. In *Inquiries into Truth and Interpretation*, chapter 5, pages 65–75. Clarendon Press, second edition, 2001. <4, 9, 79>

[18] D. Davidson. Radical interpretation. In *Inquiries into Truth and Interpretation*, chapter 9, pages 125–140. Clarendon Press, second edition, 2001. <9>

[19] D. Davidson. Belief and the basis of meaning. In *Inquiries into Truth and Interpretation*, chapter 10, pages 141–154. Clarendon Press, second edition, 2001. <11>

[20] D. Davidson. Repy to foster. In *Inquiries into Truth and Interpretation*, chapter 12, pages 171–179. Clarendon Press, second edition, 2001. <13>

[21] D. Davidson. A nice derangement of epitaphs. In *Truth, Language, and History*, chapter 7, pages 89–108. Oxford University Press, second edition, 2005. <78, 80, 81>

[22] K. DeRose. Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4):913–929, December 1992. <25>

[23] J. Fodor and E. Lepore. *Holism: A Shopper's Guide*. Blackwell, 1992. <16>

[24] R. Jackendoff. *The Architecture of the Language Faculty*. MIT Press, 1997. <52>

[25] D. Kaplan. Demonstratives. In J. Almong, J. Perry, and H. Wettstein, editors, *Themes from Kaplan*, pages 481–564. Oxford University Press, 1989. <22, 25, 26, 28, 65, 70>

[26] D. Kaplan. Afterthoughts. In J. Almong, J. Perry, and H. Wettstein, editors, *Themes from Kaplan*, pages 565–614. Oxford University Press, 1989. <26>

[27] K. Korta and J. Perry. Radical minimalism, moderate contextualism. In *Context-Sensitivity and Semantic Minimalism*, chapter 3, pages 94–111. Oxford University Press, 2007. <25, 26>

[28] E. Lepore and K. Ludwig. *Donald Davidson: Meaning, Truth, Language, and Reality*. Oxford University Press, 2005. <14, 78, 80, 81, 82>

[29] E. Lepore and K. Ludwig. *Donald Davidson's Truth-Theoretic Semantics*. Oxford University Press, 2007. <18>

[30] S.-J. Leslie. Moderately sensitive semantics. In G. Preyer and G. Peter, editors, *Context-Sensitivity and Semantic Minimalism*, chapter 5, pages 133–168. Oxford University Press, 2007. <40>

[31] S. C. Levinson. *Space in Language and Cognition: Explorations in Cognitive Diversity*. Cambridge University Press, 2003. <19>

[32] R. Montague. English as a formal language. In R. Thomason, editor, *Formal philosophy: Selected Papers of Richard Montague*. Yale Universite Press, 1974. <10>

[33] R. Montague. Universal grammar. In R. Thomason, editor, *Formal philosophy: Selected Papers of Richard Montague*. Yale Universite Press, 1974. <10>

[34] G. Nunberg. Transfers of meaning. *Journal of Semantics*, 12(2):109–132, 1995. <30>

[35] P. Pagin and F. J. Pelletier. Content, context, and compositionality. In G. Preyer and G. Peter, editors, *Context-sensitivity and Semantic Minimalism*, chapter 1, pages 25–62. Oxford University Press, 2007. <2, 35, 41, 43>

[36] J. Perry. Thought without representation. In *The Problem of the Essential Indexical and Other Essays*. Oxford University Press, 1993. <24>

[37] J. Perry. Indexicals and demonstratives. In B. Hale and C. Wright, editors, *A Companion to the Philosophy of Language*, pages 586–612. Blackwell, 1997. <25>

[38] J. Perry. *Reference and Reflexivity*. CSLI Publications, 2001. <25, 26>

[39] S. Predelli. Utterance, interpretation and the logic of indexicals. *Mind and Language*, 13(3): 400–414, September 1998. <26>

[40] S. Predelli. 'i am not here now'. *Analysis*, 58(2):107–115, April 1998. <26>

[41] F. Recanati. Unarticulated constituents. *Linguistics and Philosophy*, 25:299–345, 2002. <31>

[42] F. Recanati. *Literal Meaning: The Very Idea*. Cambridge University Press, 2004. <24, 29, 31, 33, 72>

[43] F. Recanati. Literalism and contextualism: Some varieties. In G. Preyer and G. Peter, editors, *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, chapter 7, pages 171–196. Oxford University Press, 2005. <24, 25, 26, 27, 36, 46, 50, 52, 54, 70>

[44] F. Recanati. Crazy minimalism. *Mind and Language*, 21(1):21–30, 2006. <35>

[45] F. Recanati. It's raining (somewhere). *Linguistics and Philosophy*, 30:123–146, 2007. <24, 25, 31, 68>

[46] I. Sag. Formal semantics and extralinguistic context. In P. Cole, editor, *Radical Pragmatics*, pages 273–94. London Academic Press, 1981. <52>

[47] J. Searle. Literal meaning. *Erkenntnis*, 13:207–224, 1978. <37>

[48] S. Soames. *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford University Press, 2002. <29>

[49] J. Stanley. Context and logical form. *Linguistics and Philosophy*, 23(4):391–434, 2000. <24>

[50] J. Stanley. Semantics in context. In G. Preyer and G. Peter, editors, *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, pages 221–253. Oxford University Press, 2005. <48>

[51] A. Tarski. The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research*, 4:341–375, 1944. <8>

[52] A. Tarski. The concept of truth in formalized languages. In J. Corcoran, editor, *Logic, Semantics, Metamathematics*, pages 152–278. Hackett Publishing Company, second edition edition, 1983. <4, 5, 8, 78>

[53] K. A. Taylor. A little sensitivity goes a long way. In G. Preyer and G. Peter, editors, *Context-Sensitivity and Semantic Minimalism*, chapter 2, pages 63–93. Oxford University Press, 2007. <25, 74>

[54] C. Travis. Insensitive semantics. *Mind and Language*, 21(1):39–49, 2006. <32>