

**Truth in sentential structure: A critical analysis of  
standard semantic accounts of context sensitivity**

**Alun Davies  
PhD Thesis**

**University of East Anglia  
School of Philosophy  
September 2013**

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with the author and that use of any information derived there from must be in accordance with current UK Copyright Law. In addition, any quotation or extract must include full attribution.

## **Abstract**

The principal aim of this thesis is to assess the view that the formal properties of sentences of natural languages encode truth-conditions. The question I pursue is whether truth-conditional semantic theories are capable of accounting for the various ways in which contextual factors contribute to the determination of the truth-conditional content of sentences. The thesis will assess and evaluate three different approaches standard contemporary truth-conditional semanticists have set forth in response to what I shall refer to as the challenge from pervasive context sensitivity, which is essentially the claim that context plays a more extensive role in the determination of content than that of fixing the semantic values of standard indexical expressions. I shall aim to show that each of the approaches mentioned fails to provide an adequate account of how the phenomenon of context sensitivity can be explained on the basis of our linguistic competence alone, and shall then explore some of the consequences of this.

Alun Davies  
UEA

## Contents

Introduction	p.1
Chapter 1: A Standard Semantic Theory	p.5
1. Theoretical foundations	p.5
1.1 The role and form of a semantic theory	p.5
1.2. Meaning and truth	p.9
1.3. Constraining the truth theory	p.11
1.4. Reference, satisfaction, and (absolute) truth	p.13
1.5. The context invariance of the axioms	p.14
1.6. Interpreting, understanding, and communicating	p.17
1.7. A systematic method of communication	p.18
2. Constructing a semantic theory	p.20
2.1. Structural descriptions	p.20
2.2. Lexical axioms	p.24
2.3. Compositional axioms	p.26
3. Semantic facts	p.28
3.1. Ambiguity	p.29
3.2. Negative facts	p.32
3.3. Semantic entailment	p.34
3.4. Context sensitivity	p.37
4. Summary	p.42
Chapter 2: Pervasive context sensitivity	p.43
1. Beyond standard indexicality	p.43
1.1. Context shifting arguments	p.44
1.2. Semantic incompleteness	p.47
1.3. Semantic expansion	p.49
2. Three responses	p.50
2.1. The moderate response	p.51
2.2. The radical response	p.53
2.3. The minimalist response	p.58
3. Summary	p.61
Chapter 3: Semantic Minimalism	p.62
1. The motivation for minimalism	p.62
1.1. The core claims of semantic minimalism	p.63

1.2.	C&L's semantic minimalism	p.65
1.3.	Borg's semantic minimalism	p.70
2.	Objections to semantic minimalism	p.74
2.1.	Objections to C&L's semantic minimalism	p.77
2.2.	Objections to Borg's semantic minimalism	p.81
3.	Summary	p.86
Chapter 4: The Aphonic Approach		p.87
1.	The argument from binding	p.88
1.1.	Context sensitivity and binding	p.88
1.2.	Pragmatic approaches	p.91
1.3.	Logical form, the binding assumption, and aphonic elements	p.94
2.	Assessing the aphonic approach	p.98
2.1.	An alternative account	p.98
2.2.	'Rain' again	p.107
3.	Summary	p.115
Chapter 5: The Unarticulated Constituent Analysis		p.116
1.	Unarticulated constituents	p.116
1.1.	Perry's proposal	p.117
1.2.	The lexical and conceptual adicity of 'rain'	p.123
1.3.	Rejecting metaphysical considerations	p.126
1.4.	Taylor's proposal	p.128
2.	Objections to parametric minimalism	p.131
2.1.	Understood participants	p.131
3.	Summary	p.134
Conclusion		p.136
Bibliography		p.140

(Word count: 51,000)

## INTRODUCTION

The principal aim of this thesis is to examine a certain conception of the connection between linguistic meaning and truth that has become dominant in contemporary semantics. The view can be stated as follows: to understand what a declarative sentence means is to be in a position to know what conditions must obtain in order for that sentence to be true at any context of utterance. On this view, the meaning of a declarative sentence *S*, determined exclusively on the basis of the formal properties of *S*, encodes enough information to enable speakers at contexts to assign complete sets of truth-conditions to utterances of *S*. Consequently, the role of a semantic theory is to provide a systematic method for interpreting the truth-conditions of sentences relative to contexts.

Standard truth-conditional approaches to semantics are thus predicated on the claim that notions of linguistic meaning and truth stand in a particularly close relation to one another. Foundational arguments for the claim will be examined in the opening chapter. The rest of the thesis, however, will be devoted to evaluating the force of certain problems that arise regarding the viability of the truth-conditional semantic framework. In particular, the thesis will critically evaluate how well truth-conditional semantic theories hold up against recent arguments to the effect that contextual factors play a much more pervasive role in the determination of content than was traditionally supposed.

Early advocates of truth-conditional semantic theories (e.g. Davidson, Lewis) typically assumed that there is only a very limited extent to which truth-conditional content is contextually determined, and so only a limited range of contextual factors a semantic theory would need to take into account. What is more, they argued that the effects context *does* have on content can be explained in terms of the stable, context-invariant semantic properties that standard semantic theories will be able to associate with the formal, lexico-syntactic properties of their target languages. In recent years, however, a great deal of effort has been devoted to examining the various ways in which speakers

exploit their knowledge of context in order to determine the content expressed by utterances of sentences. As a result, many now argue that the phenomenon of context sensitivity is much more pervasive in natural language than was initially assumed. This has presented a host of new challenges for standard semantic theorists.

The argumentative contribution of the thesis will be to assess and evaluate three different approaches standard semantic theorists have set forth of late in response to what I shall refer to as the challenge from the phenomenon of pervasive context sensitivity, which comes in the form of the claim that context plays a more extensive role in the determination of content than that of fixing the semantic values of overt context sensitive expressions; i.e. *standard indexical expressions*.

The three approaches I will be interested in here are the following: *semantic minimalism*, *the aphonic approach*, and *the unarticulated constituent analysis (UCA)*. Semantic minimalists maintain that the arguments that have been set forth in defence of pervasive context sensitivity are ill-founded, and that all genuine contextual effects on content are triggered by occurrences of standard indexical expressions situated in overt sentential structure. Aphonic theorists maintain that arguments for pervasive context sensitivity are well-founded, and seek to respond to such arguments by showing that all contextual effects on content are triggered by *aphonic elements* in sentential structure. UCA theorists, who also find arguments for pervasive context sensitivity persuasive, argue that all contextual effects on content are triggered by the intrinsic semantic properties of lexical items.

All three approaches thus adopt the view that the challenge from pervasive context sensitivity can be met without abandoning the crucial assumption that the lexico-syntactic properties of any given declarative sentence in a natural language encode a determinate set of truth-conditions. Each approach, however, adopts a very different conception of the way in which form and context interact in the determination of truth-conditional content. The question that motivates the thesis is whether any of these approaches provides an adequate account of that interaction.

Ultimately, the conclusions I will reach will be largely negative. That is, I shall show that each of the approaches mentioned fails to provide an adequate account of how the phenomenon of context sensitivity can be explained on the basis of our linguistic competence alone. My more positive agenda, however, will be to show that even if we reject the view that the purely formal features of sentences encode truth-conditions, there is still much to be gained from investigating how such features systematically guide and constrain the process of linguistic interpretation.

The thesis will proceed as follows. In chapter 1, I provide a detailed exposition of the theoretical foundations of the truth-conditional semantic framework. I then show how such theories might be constructed so as to provide a systematic method of interpretation for a natural language, and how they are able to accommodate a variety of important facts about the ways speakers of natural languages understand the meanings of sentences in those languages.

In chapter 2, I devote my attention to the problems the challenge from pervasive context sensitivity raises for truth-conditional semantics. I look at the arguments that have been set forth to motivate the claim that context sensitivity is a pervasive feature of natural language, and draw out the consequences these arguments have for our conception of what constitutes a semantic theory. After that, I give an overview of three different responses to arguments for pervasive context sensitivity, the aim of this being to show what position the approaches that will take centre stage throughout the remainder of the thesis occupy in the relevant theoretical landscape.

In chapter 3, I turn to two versions of semantic minimalism set forth by C&L (2005), and Borg (2004, 2012). I begin by situating semantic minimalism within a variety of different responses to the challenges raised by arguments for pervasive context sensitivity. I then provide a detailed exposition of the kinds of arguments advocates of two competing versions of semantic minimalism set forth in defence of the approach. Finally, I raise various questions about the viability of semantic minimalism. These questions focus in part on certain theory internal issues each version of the approach faces, but will also point to various problems for the claim that to acknowledge the

challenge from context sensitivity is effectively to abandon the prospects of providing any kind of systematic account of semantic competence.

In chapter 4, I assess the aphonic approach defended by Stanley (2007). My main focus here will be to assess and evaluate the arguments Stanley deploys in order to defend the claim that all genuine contextual effects on content can be traced back to aphonic elements situated in the logical forms of sentences. In evaluating the approach, I show that various counterexamples can be provided for those arguments, and seek to offer an alternative analysis of some of the examples upon which the arguments rely. I then go on to raise some potential problems for the very claim that context sensitive expressions are always accompanied by aphonic elements in sentential structure.

In chapter 5, I go on to assess the UCA approach advocated by Perry (e.g. 1986, 1998, 2011). I provide a thorough overview of the approach, arguing that it does not, as has sometimes been suggested, reject the claim that the kinds of processes speakers implement in order to incorporate contextual information into the determination of truth-conditional content are subject to lexico-syntactic constraints. Indeed, the UCA theorist takes it that all contextual effects on content are triggered and constrained by certain subsyntactic semantic properties of lexical items. Moving on, I then raise some significant problems for Perry's version of UCA, and go on to assess a modified version of his approach – *viz. parametric minimalism* – set forth by Taylor (e.g. 2001, 2007). The concern I have with Taylor's approach is that it does not overcome the kinds of problems that he himself raises for Perry's earlier version of UCA.

Finally, in conclusion, I draw together the outcomes of each chapter, with the aim of exploring some of the consequences of the arguments I set forth in each chapter. Ultimately, I shall argue that the conclusions reached in each chapter should lead us to question whether we need to commit to the kinds of core assumptions embodied in the framework of truth-conditional semantics.



## CHAPTER 1

### A STANDARD SEMANTIC THEORY

The aim of this chapter is to provide an assessment of the truth-conditional semantic framework. I begin by setting forth a detailed exposition of the theoretical foundations for truth-conditional semantic theories as developed in the work of Davidson. I then show how construction of such a theory might proceed, before setting out the key demands that a truth-conditional semantic theory will have to satisfy if it is to be considered adequate. This largely exegetical undertaking will provide a technical and conceptual framework for the more detailed analysis of contemporary semantic approaches aimed at defending the truth-conditional semantic framework that follows in subsequent chapters.

#### 1. Theoretical foundations

This section gives an overview of the theoretical foundations for standard truth-conditional approaches in semantics. The aim is to chart how the development of a truth-conditional semantic theory proceeds through the work of Davidson. Emphasis is placed on explaining: (i) what the role of a truth-conditional semantic theory is; (ii) what form such theories might take; (iii) why appealing to the notion of truth helps elucidate the notion of linguistic meaning; (iv) what kind of evidence we might call on to evaluate a semantic theory; and (v) what part such theories play in communication.

##### 1.1. The role and form of a semantic theory

Davidson (e.g. 1967, 1973) suggested that the role of a semantic theory is to provide an explicit account of what one would need to know in order to understand a language.<sup>1</sup>

---

<sup>1</sup> Davidson insisted that the kind of semantic theory he advanced need not be a theory speakers know or even believe. That is, such a theory is not conceived of as a theory of whatever mechanism speakers actually implement to acquire linguistic competence. Rather, the purpose of laying out a semantic theory, in Davidson's sense, is to articulate

What we will require from a semantic theory, on this view, is that it state information about its object language *L* in a way that would enable anyone who knew the theory to understand speakers of *L*. This imposes conditions on the form such a theory will take. First, in order for finite beings such as ourselves to be able to employ a semantic theory so as to come to understand the meanings of all the expressions in a language, the theory would have to be finitely specifiable; i.e. it would consist of a finite set of statements knowledge of which would imbue one with the practical ability to speak a language.<sup>2</sup> Second, to account for the *systematicity* and *productivity* of linguistic understanding, the theory must give a structurally revealing account of how the meanings of sentences are determined by the meanings of their constituent expressions, which is to say that a semantic theory will be compositional:<sup>3</sup>

When we can regard the meaning of each sentence as a function of a finite number of features of the sentence, we have an insight not only into what there is to be learned; we also understand how an infinite aptitude can be encompassed by finite accomplishments. (Davidson 1965, pp.8-9)

In order to show how the meanings of an unbounded number of sentences can be generated from the meanings of a finite stock of primitive expressions, it seems plausible to suppose that construction of a semantic theory will begin with the assignment of meanings to primitive expressions. One way to do this would be to identify the meanings of expressions with different entities – e.g. individuals, properties,

---

a theory knowledge of which *would* enable one to speak and understand a language *were* one to possess knowledge of it (see Davidson 1973, p.125).

<sup>2</sup> The finiteness requirement is closely related to the claim that a language must be, in principle, learnable; see Davidson 1965.

<sup>3</sup> A language counts as *systematic* if the ability to understand one of its sentences (e.g. of the form *aRb*) implies an ability to understand various related sentences (e.g. of the form *bRa*). A language counts as *productive* if its speakers are able to produce and understand an unbounded number of novel sentences.

relations, etc. – to which the expressions refer. Davidson, however, rejects this approach. On his view, it is not obvious that identifying meanings with entities will play any useful role in the construction of a semantic theory. First, assigning entities to expressions as their meanings does not, in and of itself, help explain how new meanings are generated when the meanings of expressions combine. Assigning a relation (say, an *instantiation* relation) to the syntactic operation of concatenation is of no help here, as it simply introduces another entity whose contribution to the generation of a new meaning itself stands in need of explanation.<sup>4</sup> Second, as Davidson points out, all we are really looking for when we come to the task of assigning meanings to primitive expressions is a way of specifying the effect each such expression has on the meanings of sentences in which it occurs; but once the semantic contribution of an expression has been specified, further appeal to *an entity* whose role in the theory is to do duty as the expression's meaning is superfluous (Davidson 1967, p.20). Davidson thus opts to dispense with talk

---

<sup>4</sup> To clarify, let us say we want to explain why combining the meaning of 'Mary' (call this *m*) with the meaning of 'fell' (call this *F*) gives us the meaning of 'Mary fell' (call this *Fm*) as outcome. Simply listing the meanings of the expressions  $\{m, F\}$  does not help explain why the sentence itself has the unified meaning it has. To add that there is an instantiation relation that obtains between *m* and *F* is just to add another element (call it *R*) to the list, giving us  $\{m, R, F\}$ . As a result: we enter into a vicious regress, as more and more entities are posited in the hopes of cementing the elements of the list together to form a unified meaning. As Davidson notes, Frege attempted to avoid the regress by treating the entities assigned to, say, predicates as *unsaturated*; i.e. as open functions in need of arguments. The combination of a name and a predicate is thus taken to result in saturation (thus Frege gave us the idea of characterizing the generation of the meanings of complex expressions as proceeding by way of function application between their primitive parts). Davidson, however, claims that Frege's talk of saturation merely *labels the problem* as opposed to solving it (Davidson 1967, p.17). (The problem of Bradley's regress, and the associated problem of the unity of the proposition have received a great deal of renewed attention of late: see Linsky 1992; Speaks 2008; Gaskin 2009; Collins 2012.)

of meanings as entities to which primitive expressions refer, replacing it instead with talk of the systematic contributions expressions make to the meanings of their host sentences.

Having eliminated the appeal to discrete entities to serve as meanings of primitive expressions, Davidson then turns his attention to the meanings of sentences. Let us adopt Davidson's suggestion that we can give the meaning of each expression in terms of the contribution it makes to the meanings of the sentences in which it occurs. We will then still want the theory to imply, for every sentence in the object language, an instance of the following schema:

(M)  $S$  means  $m$

where  $S$  is replaced by a (structural description of a) sentence in the object language and  $m$  is replaced by a singular term that refers to the meaning of  $S$ .<sup>5</sup> Instances of (M) thus pair sentences in the object language with their meanings, where the latter are construed as entities to which the former refer. Issues arise, though, when we come to consider what kind of expression might be used to replace  $m$  in any given instance of the schema. The problem is that simply replacing  $m$  with a singular term – e.g. ‘the meaning of  $S$ ’ – that refers to the meaning of the object language sentence will give vacuous instances of (M) – e.g. ‘Snow is white’ means ‘the meaning of *snow is white*’.<sup>6</sup> To remedy this,

---

<sup>5</sup> It is important to bear in mind that  $S$  is replaced by a *structural description of a* sentence in the object language. In order to properly capture the relevant structure of a sentence, the pertinent descriptions will sometimes be quite complex, but for ease of exposition, it is useful to use quotation mark names for sentences to replace  $S$ . Indeed, to ease exposition further still, I sometimes speak of  $S$  as being replaced by a sentence in the object language.

<sup>6</sup> Davidson (1967, p.20) describes the problem as follows: ‘Ask for example for the meaning of ‘Theaetetus flies’. A Fregean answer might go something like this: given the meanings of ‘Theaetetus’ as argument, the meaning of ‘flies’ yields the meaning of ‘Theaetetus flies’ as value. The vacuity of the answer is obvious. We wanted to know

Davidson (1967, p.22) suggests that rather than making use of a singular term that picks out a meaning, we can instead replace *m* with a use of a sentence in the metalanguage (the language we use to state the theory) that in some way *gives the meaning of*, or *translates*, the sentence in the object language. To make this explicit, we substitute (M\*) for (M):

(M\*) *S* means that *p*

where *S* is once again replaced by (a structural description of) a sentence in the object language and *p* is replaced by a sentence in the metalanguage that gives its meaning. This is a significant modification to the previous schema, as here we begin to see how a semantic theory might do its work by relating sentences in a (potentially unfamiliar) object language not to meanings as such, but to (familiar) sentences in the metalanguage whose meanings have already been grasped. A semantic theory thus comes to take the form of a theory of interpretation for an unknown language couched in a known one (Davidson 1973, p.130). The problem with (M\*), however, is that making sense of the ‘means that’ relation the schema employs involves explaining what constitutes synonymy between sentences; but the notion of synonymy raises ‘problems as hard as, or at least identical with, the problems our theory is out to solve’ (Davidson 1967, p.22). The difficulty we arrive at is that we seem to want a semantic theory that gives the meanings of sentences *without* appealing to the notion of meaning.

## 1.2. Meaning and truth

Davidson’s ‘radical’ response to the difficulties caused by appealing to meanings as is to employ a different technique for ensuring that sentences in the object language get paired with sentences that give their meanings, or say the same thing, in the metalanguage. The strategy he adopts is to replace the non-extensional ‘means that’ as

---

what the meaning of ‘Theaetetus flies’ is; it is no progress to be told that it is the meaning of ‘Theaetetus flies’. This much we knew before any theory was in sight.’

filler for the gap between the sentences with an extensional counterpart ‘is true iff’.<sup>7</sup> The resulting schema eliminates the appeal to the essentially intensional notion of *meaning* by replacing it with the purely extensional notion of *truth*:

(T)  $S$  is true iff  $p$

What we end up with is Tarski’s T-schema, which he employed in his attempt to provide a definition of the truth-predicate for formal languages. Tarski (1986) showed that in order to give such a definition, what we require is a theory that would generate, for every sentence  $S$  in an object language  $L$ , a statement of the conditions that must obtain for  $S$  to be true; i.e. for  $S$  to fall under the extension of the predicate ‘is true’. According to Tarski’s Convention T, such a definition would be *materially adequate* if for each sentence  $S$  in  $L$  the theory entailed an instance of (T) – i.e. a T-sentence – in which, as has already been implied,  $S$  is replaced by a structural description of a sentence in the object language and  $p$  is replaced by a sentence in the metalanguage that translates the sentence described by  $S$  – e.g. ‘Snow is white’ is true iff *snow is white*.

Significantly, the only semantically loaded notions the theory needs to appeal to in order to do this are the *extensional* notions of reference and satisfaction relative to infinite sequences of objects. The axioms of the theory state the satisfaction conditions of complex sentences by showing how those conditions derive from the satisfaction conditions of simple open sentences, and state the satisfaction conditions for open sentences relative to sequences.<sup>8</sup> Truth is then defined as a special case of satisfaction: a

---

<sup>7</sup> Davidson describes the relevant shift as follows: ‘...let us try treating the position occupied by  $p$  extensionally: to implement this, sweep away the obscure ‘means that’, provide the sentence that replaces  $p$  with a proper sentential connective, and supply the description that replaces ‘ $s$ ’ with its own predicate’ (Davidson 1967, p.23).

<sup>8</sup> For example, the satisfaction conditions for the open sentence ‘ $x$  swims’ are given by assigning an index  $i$  to the variable  $x$  that relates it to a particular position in any given sequence of objects. A sequence then satisfies the open sentence iff its  $i$ ’th member swims. The relevant axiom might thus read as follows:  $\text{SAT}[\Sigma, x_i \text{ swims}]$  iff  $\sigma_i$  swims,

true sentence is one satisfied by all sequences; a false sentence one satisfied by none. Davidson's insight was that since *giving the truth-conditions* of a sentence is one way of *giving its meaning*, a Tarski-style truth theory that issued in a T-sentence statement of the truth-conditions for every sentence in the object language would, *under certain constraints*, do duty as a semantic theory.<sup>9</sup> All reference to entities (e.g. *facts*) that serve as the meanings of sentences thus drops out, which means meanings have been altogether eliminated from the theory.

### 1.3. Constraining the truth theory

I mentioned that in order for a truth theory to serve as a semantic theory, it must conform to certain constraints. Why would we be required to impose constraints on a truth-theory? What constraints need to be imposed? To answer these questions, let us begin by reminding ourselves that Tarski's aim is to define truth, *not* to gain purchase on

---

where  $\sigma_i$  is the  $i$ 'th member of the sequence,  $\Sigma$ . Satisfaction conditions can also be assigned to referring expressions (e.g. names), but we will tend to speak of axioms for referring expressions as stating *reference conditions*, which simply state what individual a referring expression refers to; e.g. 'Bill' refers to Bill. In the case of satisfaction conditions for sentences, we say that a sequence satisfies, say, 'Bill swims' iff the member of the sequence referred to by 'Bill' swims. Of course, if a sentence is satisfied by one sequence then it is satisfied by all; which means the sentence is true; which means truth is a special case of satisfaction.

<sup>9</sup> The idea of specifying the meaning of a sentence in terms of its truth-conditions goes back at least as far as Frege. Davidson's approach can be regarded as a way of implementing Frege's insights regarding the relation between meaning and truth in such a way as to construct a workable semantic theory. Other truth-conditional approaches – such as that developed by Lewis (1970) whose work builds on the work of Montague (e.g. 1973) – do the same thing in a different way. In particular, theories in the Motogovian tradition tend to be less wary of specifying the meanings of expressions by way of an appeal to meanings as intensional entities.

the concept of meaning. Consequently, in stating his Convention T, Tarski gets to help himself to the assumption that the sentence on the right-hand side of the biconditional in any given T-sentence will be a translation of, or mean the same as, the sentence described on the left-hand side. Let us refer to any such T-sentence as an *interpretive* T-sentence. The problem Davidson faces is that since his aim is to gain purchase on the concept of meaning, he is not able to help himself to the pertinent assumption; i.e. he is not able to treat *meaning* as primitive.<sup>10</sup>

What is required, then, is that the truth theory be constrained in such a way as to ensure, without any appeal to semantically loaded notions (such as that of *synonymy*), that the T-sentences the theory generates are interpretive. To achieve this, Davidson (e.g. 1967) begins by arguing that the totality of T-sentences generated by the theory must be true. In and of itself, however, this is not enough of a constraint to ensure that the T-sentences will be interpretive (Davidson 1973, p.134). We must thus place a further constraint on the theory by demanding that its T-sentences accurately reflect attitudes speakers of the object language have as to when sentences of the object language would hold true.

To show how the constraint operates, we start by adopting the *principle of charity*, which is to say we assume that speakers of the object language have accurate beliefs about their environment, and that they take the sentences they utter to hold true relative to those beliefs. Accordingly, when a T-sentence generated by our theory tells us that a sentence in the object language is true iff, say, there is snow on the ground, it follows

---

<sup>10</sup> The assumption is fine as long as the object language is contained in the metalanguage. In such a case, the interpretive character of the T-sentences generated by the theory is guaranteed by the fact that the very same sentence described on the left is used to replace the sentence on the right. Davidson, however, wants a theory that relates sentences in a potentially unknown object language to interpretations of those sentences in the known metalanguage. That is, he wants the theory to be one that might be of use to the radical interpreter whose only evidence for testing the consistency of his theory, as we shall see, is evidence about when speakers of the object language hold their sentences true (Davidson 1973).



that, *ceteris paribus*, there will be snow on the ground when a speaker of the object language utters that sentence. If things are not that way, then given the assumptions embodied in the principle of charity, we have evidence that the theory has assigned the wrong truth-conditions to the pertinent sentence. Of course, in many cases, identifying the relevant evidence will be more difficult. The point, though, is that when situations such as the one just described arise, the theorist, conceived of as someone engaged in the project of radical interpretation, will be required to revise the axioms of the theory in such a way as to ensure that its T-sentences fit with evidence regarding the holding-true attitudes of speakers of the object language (whatever form such evidence takes).

To sum up, we have seen: (i) how a Tarski-style truth theory provides us with the machinery for generating T-sentences; and (ii) what constraints we might impose on such a theory in order to ensure that the T-sentences it generates are interpretive. Bringing all this together, what we obtain is a theory that answers to a set of formal and empirical requirements, which, taken together, conspire to ensure that its theorems provide us with interpretations of, and so a means of understanding, the endlessly many sentences in the object language to which the theory is applied. That is, we end up with an outline of a semantic theory whose construction recovers ‘the structure of a very complicated ability – the ability to speak and understand a language’ (Davidson 1967, p.25).

#### 1.4. Reference, satisfaction, and (absolute) truth

At this point, it will be useful to consider in more detail what role is played by the extensional notions of *reference*, *satisfaction*, and *truth* in the construction of a semantic theory.<sup>11</sup> With regards truth, it is vital to keep in mind that when we come to speak of the truth-conditions of sentences, we are employing an *absolute* notion of truth; i.e. we are speaking of sentences as *true relative to the world*, rather than some *model*, *interpretation*, or *set of possible worlds* (Davidson 1977, p.216). That Davidson himself

---

<sup>11</sup> This description of the notions of reference, satisfaction, and truth as the ‘genuinely semantic notions’ echoes Lewis 1970, p.19.

adheres to an absolute notion of truth is evident from the above remarks on his position, where it was made clear that he intends the adequacy of the theory to be tested on the basis of evidence regarding whether the T-sentences it generates succeed in relating sentences in the object language to beliefs speakers of the object language have *about the way the world must be* for those sentences to hold true. The adequacy of the theory as a whole thus hangs on the question of whether its T-sentences accurately reflect the way speakers put their sentences to use when they come to talk about the world, to communicate their beliefs with one another, and so on.

Now, just as T-sentences must relate sentences to the world in the right sort of way, lexical axioms that state reference and satisfaction conditions for primitive expressions must also, in some sense, relate those expressions to the world in the right sort of way. That said, though, it is crucial to observe that Davidson never intended that the axioms of the theory should be thought of as answerable to direct evidence regarding what kind of things speakers *actually* talk about when they come to put their expressions to use. Rather, the adequacy of the axioms is to be evaluated only *indirectly*, relative to the adequacy of T-sentences they serve to generate. This leads Davidson to think of the notions of reference and satisfaction as playing a purely *instrumental* role within the context of the theory as whole; i.e., to treat the relevant notions as ‘theoretical constructs whose function is exhausted in stating the truth conditions for sentences’ (Davidson 1977, p. 223). Consequently, when we talk of words as referring to individuals, or as satisfied by sequences, we are not describing relations that actually obtain between words and entities in the world. Rather, the motivation for such talk is purely theoretical, to be dispensed with once truth-conditions have been assigned to sentences in the object language.

#### 1.5. The context-invariance of the axioms

Keeping the points made in the previous subsection in mind, note that even if we treat the contributions axioms make to the theory as purely instrumental, there are still things to be said about the way those axioms operate. Perhaps the most important feature of the axioms of a semantic theory is that they must provide us with *context-invariant* specifications of the contributions expressions make to the truth-conditions of sentences

in which they occur as well as the semantic effects of combining those expressions. Such context-invariance can be thought of in two ways, depending on the sense in which we use the word ‘context’. First, there is *linguistic (or sentential) context*, which consists of the linguistic environment in which an expression occurs.<sup>12</sup> Second, there is *extra-linguistic (or utterance) context*, which consists of a set of pragmatic factors that are in play when a sentence is uttered. A lexical axiom can be regarded as context-invariant in the first sense just if it specifies, in univocal terms, the semantic contribution an expression makes across all linguistic contexts, and in the second sense just if it specifies, in univocal terms, the semantic contribution an expression makes across all extra-linguistic contexts of utterance. In order for the axioms to serve their purpose in a semantic theory of the kind described, they must be context-invariant in both senses.<sup>13</sup>

The reason for this is that an essential feature of the kind of semantic theory described so far is that its axioms must form a finite base from which T-sentence statements of the truth-conditions of sentences in the language can be recursively derived. But there is an unbounded number of novel linguistic contexts in which any given expression might occur, and an unbounded number of novel extra-linguistic contexts in which sentences containing the given expression might be put to use. So if we attempted to specify the semantic contribution of each expression relative to each particular context in which the

---

<sup>12</sup> So for example, in the following sentences, ‘break’ appears in different linguistic contexts: (a) Bill broke the vase; (b) The vase broke; (c) Mary broke her arm; (d) The break was quite bad; (e) The broken vase had been priceless; (f) There goes another broken heart. In some of the cases listed the linguistic context varies only to the extent that the verb ‘break’ combines with different subjects and objects. In others, the linguistic context has an effect on the expression’s syntactic role: i.e. in a, b, and c, it appears as a verb; in d it appears as a noun phrase; in e and f it appears as an adjective.

<sup>13</sup> I focus here specifically on lexical axioms because they will be of the kinds of axioms that interest us most in what follows. Compositional axioms are also context-invariant, but stating what invariance amounts to relative to such axioms is somewhat different from stating what it amounts to for lexical axioms.

expression is apt to occur, our ‘theory’ would no longer be one that had the form of a Tarski-style truth theory.

The point can be further illustrated by thinking about the task of the radical interpreter. As we saw, the radical interpreter is engaged in the ongoing task of testing the adequacy of her T-sentences in light of evidence gleaned from particular occasions on which speakers of the object language put their sentences to use. When a T-sentence is shown not to reflect the holding-true attitudes of speakers relative to some such occasion, the theorist then modifies some of the axioms so as to come up with a more fitting theory. In effect, her task is to form generalisations about the contributions expressions make that will enable her to predict the truth-conditions of sentences on subsequent occasions of their use. But if the radical interpreter were to assume that each expression makes a different semantic contribution at each novel linguistic or extra-linguistic context, she would not be able to form the relevant generalisations. It is thus hard to imagine how she would ever advance her understanding of the language.

A workable truth-conditional semantic theory, then, is one that operates with a fixed, finite set of context-invariant axioms, which specify the semantic contributions expressions make across all contexts. A theory that did not avail itself of a finite set of context-invariant axioms would not have the form of a truth theory. It would not be of use to the radical interpreter whose aim is to ensure that its T-sentences conform to the evidence. And it would not account for the systematic and productive manner in which speakers are able to pair sentences with truth-conditions.<sup>14</sup>

---

<sup>14</sup> This may look like an exceedingly obvious point, not worthy of nearly so much labour. However, given that much of the following thesis will be concerned with the question of the extent to which extra-linguistic context contributes to the determination of truth-conditional content, it is worthwhile dwelling on such issues. For the time being, I shall hold off saying anything about how the axioms of a truth-conditional semantic theory might accommodate extra-linguistic contextual effects on content without giving up the invariance of its axioms. We will come back to this.

## 1.6. Interpreting, understanding, and communicating

We have so far been taking it that a semantic theory is something a theorist employs in order to interpret an object language. Looking at things in this way, we take it that for each object language, there is *a* semantic theory that explicitly articulates the linguistic competence of all speakers of that language. Put another way, we can say that all speakers of an object language are equipped with the same semantic theory for their language.<sup>15</sup>

The view makes a great deal of sense, especially when we come to think about how speakers are able to routinely engage in successful communication with each other. Assuming that speakers who share the same language enter into communicative exchanges equipped with the same semantic theory for interpreting the sentences uttered in such exchanges, it is no surprise that they routinely succeed in the task of understanding each other. The theory they employ can thus be thought of as *the* semantic theory for *the* object language they share.

That said, though, on the basis of observations regarding ways in which speakers are able to break with normal linguistic conventions during communication, Davidson (1983) came to question whether successful communication would require that all speakers enter into communicative exchanges equipped with the same semantic theory. He suggested that whilst a speaker may well be equipped with a *prior theory* that provides her with a ‘systematic method’ for interpreting the sentences uttered by her interlocutors, she will sometimes be required modify that theory to accommodate (what are from their perspective) unusual uses of expressions. Such modifications yield a *passing theory* that is geared toward producing accurate interpretations of sentences uttered by her interlocutors on that occasion. On this view, then, successful

---

<sup>15</sup> Talking like this is fine as long as we keep in mind that speakers are only ‘equipped with the theory’ in the sense that they are equipped with a competence that is explicitly articulated by the theory. Remember: the theory merely expresses the structure of the linguistic competence of speakers – it is not necessarily something speakers *know*.

interpretation requires convergence, not of prior, but of passing theories. In some cases, for various reasons, modifications made to a prior theory in order to produce an effective passing theory on a given occasion might end up being incorporated into a speaker's prior theory. In others, the modifications made are taken to serving a purpose only relative to a particular occasion, and are thus dispensed with once the relevant exchange terminates.

The important point here is that once we allow that success in interpreting the utterances of other speakers need only depend on convergence of passing theories, the idea that speakers who are able to communicate must all enter into acts of communication equipped with the same semantic theory drops out. Davidson's later work thus marks a significant shift from his earlier views. The very idea that there are object languages, conceived of as publically available languages governed by a fixed semantic theory, is no longer essential. Interpretation does not target object languages, but the idiolects of individual speakers. Indeed, each communicative exchange we enter into comes to be thought of as involving radical interpretation of an idiolect. In fact, Davidson (1983, p.265) goes so far as to say that there are no such things as shared, publically available, object languages.<sup>16</sup>

#### 1.7. A systematic method of interpretation

Now that we pieced together a description of how a truth-conditional semantic theory of the kind set forth by Davidson operates, we come to the view that the role of such a theory is to provide a systematic method of interpretation.<sup>17</sup> Threading together the

---

<sup>16</sup> The idea that speakers share a language, English say, can now be thought of in terms of speakers having idiolects that overlap with each other to a considerable extent. Their prior theories will thus be very close, and communication will, in a significant number of cases, proceed without the need for modification of those theories.

<sup>17</sup> As has already been mentioned (see p.6, n.3), *systematicity* is that feature of linguistic understanding that ensures that being in a position to understand one occurrence of an expression will ensure the ability to understand other occurrences of that expression. If

various elements of the theory, we end up with the following overview of the way the method of interpretation delivered by a semantic theory works:

- (i) The objects of interpretation for a semantic theory are sentences belonging to the idiolects of speakers.
- (ii) To yield interpretations of such sentences, we employ a Tarski-style truth theory that generates T-sentence statements of the meanings of sentences exclusively in terms of their compositionally determined truth-conditions.
- (iii) The context-invariant axioms of the theory are open to continual revision based on evidence regarding whether T-sentences generated by the theory adequately reflect the holding-true attitudes of the speaker whose idiolect the theory targets.
- (iv) The adequacy of the theory in its capacity as a method of interpretation for a target idiolect is open to empirical investigation, to be evaluated in terms of how well the T-sentences generated by the theory have thus far conformed to the relevant evidence.

A crucial assumption that underlies the truth-conditional conception of what constitutes a semantic theory is that such a theory can assign truth-conditions to sentences purely on the basis of their formal features. That is, to derive a T-sentence statement of the truth-conditions of a sentence *S*, we need only look to the lexical axioms for expressions in *S* and the compositional axioms relevant to the syntactic arrangement of *S*. No other factors need enter into interpretation in order for a complete set of truth-conditions to be assigned to any given sentence uttered. The formal features of sentences, in short, encode truth-conditions.

---

there were not some stable, systematic relation between diverse occurrences of expressions, it would be hard to explain how anyone could understand a language. A semantic theory, conceived of as a method of interpretation, is thus systematic if it establishes stable relations between different occurrences of expressions. The theory we have described thus far inherits this feature on the basis of its having a fixed, finite set of context-invariant lexical axioms.

That said, though, a significant problem that looms on the horizon for our truth-conditional semantic theory is that there are cases where extra-linguistic contextual factors seem to play an indispensable role in determining the truth-conditions of sentences. But if this is so, it appears that sentences themselves do not encode complete sets of truth-conditions after all. We are thus left questioning whether a semantic theory of the kind described thus far will be able to account for the ways in which extra-linguistic contextual factors seem apt to contribute to the determination of truth-conditional content. This is the question that will motivate much of the following thesis.

However, before we go on to pursue matters of context any further, the next section will be devoted to giving a more concrete account of how we might go about the task of constructing a semantic theory for a language, focussing in particular on the formal methods that might be adopted in order to achieve this.

## 2. Constructing a semantic theory

Having gone to some lengths to describe the theoretical foundations for truth-conditional approaches in semantics, I now want to deal with the practical question of how we might go about the task of constructing a truth-conditional semantic theory. In doing so, I shall take it that the object of interpretation for the theory is a natural language (setting aside the question of whether natural languages are to be characterised in terms of public languages, or the idiolects of speakers). That said, there are three kinds of information a semantic theory must provide us with in order to fulfil its role as a systematic method of interpretation for sentences: (i) information about the semantically relevant structure of each sentence in the target language; (ii) information about what contribution each expression makes to the truth-conditions of sentences in which it occurs; and (iii) information about the kinds of semantic effects that result from combining expressions as per each sentence's syntactic arrangement.

### 2.1. Structural descriptions

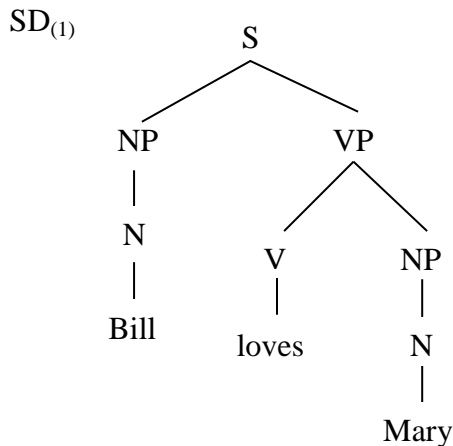
The meaning of a sentence is partly determined by its structure. That said, in order to give an account of what each sentence means, one of the things we need to get clear on is how the constituent expressions of sentences are structurally related to one another.



To that end, we need to adopt some method of describing the underlying syntactic arrangement of the strings of signs or sounds that constitute the surface forms of sentences. For example, consider (1):

(1) Bill loves Mary

The following (somewhat simplified) *phrase marker* provides a *structural description* (SD) of the way the constituent expressions of (1) are combined with one another:<sup>18</sup>



---

<sup>18</sup> Phrase markers, given here in the form of a tree diagram, represent relations of dominance and precedence over sets of labelled nodes. Since we shall be using phrase markers of this kind to give structural descriptions of sentences, it will be useful to set out some terminology for describing the diagram: (a) each node in a tree diagram dominates all nodes below it that can be traced back to it by connecting branches; (b) the root node is the node that dominates all others, but is not dominated; (c) if node A dominates node B, and no node intervenes between them, then A *immediately dominates* B, and B is A's daughter node; (d) two nodes immediately dominated by the same node are sister nodes; (e) a node that dominates no other is a terminal node; (f) nodes with two daughter nodes are branching nodes, and nodes with just one daughter are non-branching nodes; (g) precedence is a left-to-right relation between nodes that do not stand in the dominance relation to each other. For more on the properties of trees and their applications in semantics, see Higginbotham 1985; Partee et. al. 1990; Larson and Segal 1995; Heim and Kratzer 1998.

We shall adopt the somewhat standard assumption that phrase markers such as  $SD_{(1)}$  are the outputs of *syntactic* derivations that yield well-formed sentences, and are the inputs to *semantic* derivations that yield truth-conditions (Heim and Kratzer 1998, p.45; Larson and Segal, p.104). That said, it will be useful to give an outline (albeit a very general one for now) of the layout of phrase markers and some of the principles that govern their organisation.

An important idea to begin with is that the expressions, which we shall often refer to as lexical items, which occupy the terminal nodes in the phrase marker, *project* certain features, such as their syntactic category, onto their dominating nodes. (We shall say that each lexical item has a corresponding *lexical entry*, which specifies syntactically *and* semantically relevant information about that item.) This is captured in  $SD_{(1)}$  in that the node immediately dominating each lexical item is labelled according to the item's syntactic category; i.e. N stands for 'noun'; V for 'verb'. Lexical items also project information regarding what kinds of sentential constituents they need to combine with in order to form phrases. For instance, part of the information encoded in a lexical entry specifies what *internal arguments* it selects for (if any) in order to form phrases, where the selecting item is referred to as the *head* of the resulting phrase, and the constituents that serve to fill its argument positions are its *complements*.<sup>19</sup> Simplifying somewhat, the transitive verb 'loves' takes a noun phrase (NP) complement to form a verb phrase (VP), whereas the nouns 'Bill' and 'Mary' (whose entries list no mandatory arguments) form

---

<sup>19</sup> Lexical items that fill internal argument positions are sometimes realised as subjects at surface form. This occurs in unaccusative constructions, where a deep-structure object moves to an unoccupied subject position (e.g. 'The water froze'), and passive constructions, where the external argument is suppressed, again causing the deep-structure object to move so as to ensure that the sentence has a subject at surface form (e.g. 'The bank was robbed'). According to GB-theory (Chomsky 1981), movement of complements occurs between D-structure and S-structure (two level of linguistic representation) occurs between complements of the relevant verbs undergo movement from object positions at D-structure to the subject position at S-structure.

NPs without taking complements. Verbs like ‘love’ also select an *external argument* that combines with the VP formed by the item and its internal argument.<sup>20</sup> So, in  $SD_{(1)}$ , the VP ‘loves Mary’ combines with the NP ‘Bill’ to satisfy the requirement for an external argument projected by the verbal head. What we end up with, then, is a representation of the syntactic arrangement of (1) that reveals the order in which its constituent expressions combine to form phrases, which then combine to form a complete sentence.

One other important point to note about structural descriptions is that they are taken to reflect a particular level of linguistic representation. It should already be obvious that phrase markers encode various kinds of information that are not apparent from the surface forms of sentences; e.g. they represent boundaries between the phrasal constituents of sentences and the hierarchical organisation of sentences in a way that linearly ordered surface forms do not. Thus, there seem to be good reasons for differentiating between the surface forms of sentences, which constitute one level of linguistic representation, and the structural descriptions of sentences, which represent another. Following Chomsky (2005), we can think of the computational procedures via which sentences are syntactically generated as issuing in pairings of sounds and meanings. That is, we take each sentence as a  $\langle PF, LF \rangle$  pair, where PFs are representations that encode information pertaining to the sentence’s phonological or phonetic features, and LFs are representations that encode information pertaining to the sentence’s semantic features. I shall therefore treat PFs as linearly ordered strings of signs or sounds that constitute the surface forms of sentences, and LFs as linguistic representations that encode all grammatical information that is relevant for semantic

---

<sup>20</sup> Chomsky (1965) refers to that component of a lexical entry that specifies the number of internal arguments an item takes as the item’s *subcategorization frame*. The notion of argument structure is more general in that it also specifies arguments that appear outside the phrase an item forms via combination with its complements.

interpretation. Accordingly, we will presume that structural descriptions mirror the LFs of sentences.<sup>21</sup>

## 2.2. Lexical axioms

The next step in showing how a semantic theory determines the truth-conditions of (1) is to supplement our theory with semantic axioms that specify the meanings of the simple expressions of (1) and the semantic effects of combining those expressions. Each lexical axiom, which constitutes the semantic component of a lexical entry, specifies the meaning of an expression in terms of the contribution the expression makes to the truth-conditions of all sentences in which it occurs. We shall refer to the semantic contributions expressions make to the truth-conditions of sentences as their *semantic values*.

Following Heim and Kratzer (1998), I give the semantic value of each simple expression (or lexical item) in terms of its extension: i.e. the semantic values assigned to the proper names ‘Bill’ and ‘Mary’ are the individuals to which they refer, and the semantic value assigned to the transitive verb ‘loves’ is the function from individuals to a function from individuals to truth-values that that expression denotes. This can be represented as follows (where ‘ $\| \dots \|$ ’ is the interpretation function that assigns semantic values to the enclosed expressions):

---

<sup>21</sup> For introductions to the relevant notions see Radford (1988), Haegeman (1994), and Hornstein, *et al.* (2005). It is important to emphasise that whilst LFs encode certain semantically relevant pieces of information about sentences, and are thus named for their affinity with the logical forms attributed to sentences by semanticists, the linguist’s LF representations need not necessarily encode *all* the information encoded in the semanticist’s logical forms. As Haegeman writes: ‘The LF representation is not identical to the semantic representations introduced by semanticists and formal logicians; LF is an intermediate step that mediates between S-structure and the semantic representations’ (Haegeman 1994, p.492). The issue of just how close a connection there is between a sentence’s LF and its logical form will emerge later chapters.

- (i)  $\| \text{'Bill'} \| = \text{Bill}$
- (ii)  $\| \text{'Mary'} \| = \text{Mary}$
- (iii)  $\| \text{'loves'} \| = \lambda x. \lambda y. y \text{ loves } x$ <sup>22</sup>

Our truth-conditional semantic theory will thus make use of lexical axioms that explicitly state the semantic values of expressions by stating what entities (i.e., individuals and functions) lexical items denote. This might appear to put our theory somewhat at odds with the kind of Davidson-style semantic theory described in the previous section. That said, though, the decision to represent the semantic values of lexical items in the above way is motivated by nothing other than the fact that such representations specify the contributions lexical items make to truth-conditional content in a particularly perspicuous manner. All reference to entities as the denotations of expressions should thus be thought of as purely instrumental. So when we talk of a predicate as denoting a function from individuals to truth-values, we are not thereby committing ourselves to the claim that there is some entity, i.e. the function itself, which is *the meaning of* that predicate. Rather, we are merely employing a useful method for stating what kind of contribution the expression makes to the truth-conditions of sentences in which it occurs.<sup>23</sup>

---

<sup>22</sup> The lambda operator  $\lambda$  forms expressions that refer to functions by abstracting over variables in open sentences. So, applying  $\lambda$ -abstraction to the open sentence ' $x$  smokes' (where  $x$  ranges over the domain of individuals ( $D_e$ )) yields the expression ' $\lambda x \in D_e. x$  smokes', which refers to a function that maps every individual  $x$  to *true* iff  $x$  smokes. When the function is applied to an argument of the appropriate type,  $\lambda$ -reduction applies, and the argument replaces the bound variable. See Dowty *et al.* (1981) and Gamut (1991) for more on the kind of type-theoretic apparatus employed in natural language semantics.

<sup>23</sup> It should be noted that the chosen method for specifying semantic values follows the format more commonly employed in model-theoretic approaches to semantics, which stem from the work of Montague (1973). Nevertheless, there is no reason to think that the decision to specify semantic values in this way bars us from treating lexical entries

### 2.3. Compositional axioms

Having provided those lexical axioms that we will make use of in determining the truth-conditions of (1), we now need to show how the semantic values assigned to the tree's terminal nodes project to its higher nodes. In the case of non-terminal, non-branching nodes (such as 'V' in  $SD_{(1)}$ ) we can give a simple semantic rule (NN) that states that such nodes inherit their semantic values directly from their daughter nodes:

$$(NN) \left\| \begin{array}{c} \beta \\ | \\ \alpha \end{array} \right\| = \|\alpha\|$$

In order to assign values to branching nodes, a compositional axiom is required that states the semantic effect of syntactically combining the daughter nodes. Given the path we have taken so far, it makes sense to construe the semantic counterpart of syntactic concatenation as *functional application* (FA). The following thus serves as a general compositional axiom:

$$(FA) \left\| \begin{array}{c} \gamma \\ \swarrow \quad \searrow \\ \alpha \quad \beta \end{array} \right\| = \|\beta\| (\|\alpha\|)$$

This says that for any pair of sister nodes  $\alpha$  and  $\beta$ , if the semantic value of  $\alpha$  is in the argument domain of the (functional) semantic value of  $\beta$ , the value of the immediately

---

as specifying *absolute* statements of the reference/satisfaction/truth-conditions of expressions. Indeed, (i) and (ii) should be read as stating that the expressions 'Bill' and 'Mary' refer to Bill and Mary, respectively. Similarly, (iii) can be read as the statement that 'loves' is true of any pair of individuals in the world whose first member loves the second member; or is that predicate that is *satisfied* by any such pair. None of this puts us at odds with Davidson. Indeed, theorists working toward the development of a workable Davidson-style semantic theory have adopted different ways of specifying the semantic values of expressions: see Larson and Segal 1995; Lepore and Ludwig 2007.

dominating node  $\gamma$  is the value of the function assigned to  $\beta$  for the value of  $\alpha$  as argument.<sup>24</sup>

Finally then, having assigned lexical axioms to all the constituent expressions of (1), and with a compositional axiom that states the semantic effects of syntactic concatenation, we can now derive the truth-conditions for (1) as follows (where a flattened out version of  $SD_{(1)}$  is used for sake of clarity):

$$\begin{aligned}
 & \left\| \left[ S \left[ NP \left[ N \text{ Bill} \right] \right] \left[ VP \left[ V \text{ loves} \right] \left[ NP \left[ N \text{ Mary} \right] \right] \right] \right] \right\| \\
 &= \left\| \left[ VP \left[ V \text{ loves} \right] \left[ NP \left[ N \text{ Mary} \right] \right] \right] \right\| \left( \left\| \left[ NP \left[ N \text{ Bill} \right] \right] \right\| \right) && \text{by FA} \\
 &= \left\| \left[ V \text{ loves} \right] \right\| \left( \left\| \left[ NP \left[ N \text{ Mary} \right] \right] \right\| \right) \left( \left\| \left[ NP \left[ N \text{ Bill} \right] \right] \right\| \right) && \text{by FA} \\
 &= \left\| \text{'loves'} \right\| \left( \left\| \text{'Mary'} \right\| \right) \left( \left\| \text{'Bill'} \right\| \right) && \text{by NN} \\
 &= [\lambda x. \lambda y. y \text{ loves } x] (\text{Mary}) (\text{Bill}) && \text{by (i), (ii), (iii)} \\
 &= [\lambda y. y \text{ loves Mary}] (\text{Bill}) && \text{by } \lambda\text{-reduction} \\
 &= \text{true iff Bill loves Mary} && \text{by } \lambda\text{-reduction}
 \end{aligned}$$

The derivation yields a proof of the following statement: ‘Bill loves Mary’ is true iff Bill loves Mary. Such disquotational statements of the truth-conditions of sentences reveal what is (at the very least) an important part of what speakers know in understanding what those sentences mean. As Higginbotham writes:

---

<sup>24</sup> Functional application does not always work. For example, neither of the functions assigned to constituents of the phrase ‘wild animal’ can be applied to the other. We must thus use a different compositional rule to compute a value for the phrase; *viz.*, *predicate modification*. On this rule: ‘Bill is a wild animal’ is true iff Bill belongs to the intersection of the set of wild things and the set of animals. See Higginbotham 1985, Heim and Kratzer 1998, Pietroski 2006.

...disquotational statements of truth-conditions are significant for their combination of depth and obviousness: the whole lot of them are among the things one knows when one knows a language, and they remain a touchstone of adequacy even whilst the theory is elaborated. (Higginbotham 1988, p.31)

As we saw in the previous section, such statements effectively serve to connect sentences in the language targeted by the theory with the things in the world speakers of the language use those sentences to talk about. What the derivation ultimately reveals, then, is how the truth-conditional semantic content of (1) is composed 'by a finite number of applications of some of a finite number of devices that suffice for the language as a whole, out of elements drawn from a finite stock (the vocabulary) that suffices for the language as a whole' (Davidson 1968, p.94).

Having thus demonstrated how knowledge of a fragment of a semantic theory would enable us to determine the truth-conditions of *one* sentence, we can think of the task of a complete truth-conditional semantic theory for a natural language as that of compiling an inventory of lexical and compositional axioms from which we would be able to derive equally adequate statements of the truth-conditions of *all other* sentences in that language.

### 3. Semantic facts

We have seen that a truth-conditional semantic theory is a theory whose purpose is to account for a range of facts pertaining to the ways in which speakers interpret the meanings of expressions. We have so far seen that one of the kinds of facts we can deduce from a truth-conditional semantic theory are what Higginbotham (1988) refers to as *disquotational facts*; e.g. the fact that 'Bill loves Mary' is true iff Bill loves Mary; that 'Mary' refers to Mary; that 'loves' is true of an ordered pair of individuals  $x$  and  $y$  iff  $x$  loves  $y$ , and so on. We now turn to a variety of other facts about the ways speakers understand expressions that a semantic theory must be able to explain.



### 3.1. Ambiguity

Ambiguity occurs when speakers are able to associate more than one set of truth-conditions with a given sentence. The following examples serve to demonstrate different ways in which ambiguity arises:

- (2) a. Bill went to the bank
- b. Bill called John from London
- c. Every boy loves some girl
- d. The duck is ready to eat

Adopting the approach outlined in the previous section, we can take it that ambiguity arises when two expressions (simple or complex) have distinct meanings even though they are associated with the same sign or sound, or with the same string of signs/sounds. That is, cases of *lexical ambiguity* arise when we have two expressions that mean different things even though they are both associated with the same sign. The truth-conditions of (2a) thus vary depending on whether the sign ‘bank’ is understood to stand for an expression that denotes a financial institution, or an expression that denotes a place by a river. Standard truth-conditional semantic theories have accommodated this fact by providing two distinct lexical axioms for ‘bank’: i.e.,  $\| \text{‘bank’}_1 \| = \lambda x. x \text{ is a } \textit{financial} \text{ bank}$ ;  $\| \text{‘bank’}_2 \| = \lambda x. x \text{ is a } \textit{river} \text{ bank}$ .

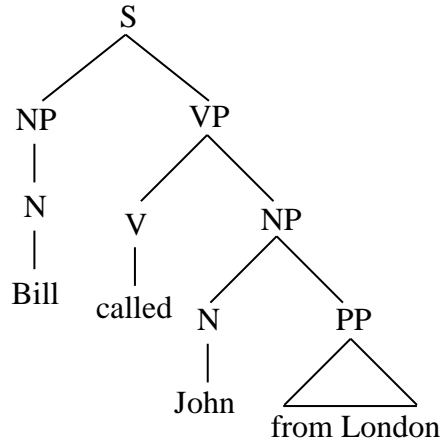
The other sentences of (2) involve some form of *structural ambiguity*, which occurs when two sentences have distinct LFs even though they share the same surface form, or PF.<sup>25</sup> So, speakers can either interpret (2b) as meaning that Bill called John, and John is from London, or that Bill called John, and the call was made from London. The difference is entirely structural, as it hangs on whether the prepositional phrase (PP)

---

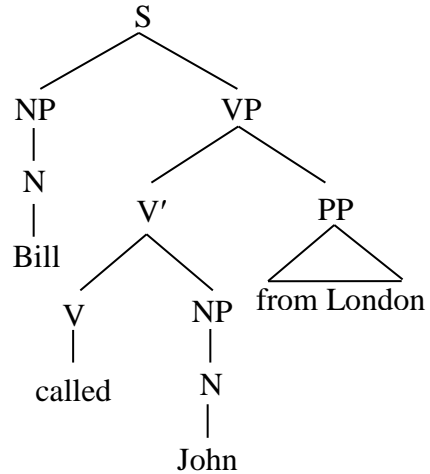
<sup>25</sup> Since we made the decision (in section 1.2.) to individuate sentences as pairings of sounds and meanings, i.e. as  $\langle \text{PF}, \text{LF} \rangle$  pairs, it would seem to make little sense to speak of *sentences* as structurally ambiguous. That said, we can take it that when we speak of *ambiguous sentences* we are simply individuating sentences, somewhat loosely, in terms of their surface forms.

‘from London’ is understood as modifying the NP ‘John’, or as modifying the VP ‘called John’. Accordingly, our semantic theory will assign two different structural descriptions to (2b):

SD1<sub>(2b)</sub>



SD2<sub>(2b)</sub>



In the case of (2c), the truth-conditions speakers are able to associate with this particular string of signs vary depending on which of the quantifier phrases takes widest scope at LF. If the quantifier denoted by the determiner phrase (DP) ‘some girl’ is taken to have wide scope, then (2c) is true iff there is some particular girl that every boy loves. Alternatively, if the quantifier ‘every boy’ takes wide scope, then we interpret (2c) as being true iff for every boy, there is some girl or other he loves.<sup>26</sup>

As indicated, differences in scope relations are represented at LF. Hence, (2c) can be assigned the following structural descriptions:

---

<sup>26</sup> This is precisely the kind of case that led Frege and Russell to argue that we should be distrustful of the idea that the superficial surface forms of sentences provide a good indication of their underlying logical forms. The predicate calculus their work gave rise to thus provides us with a means of disambiguating (2c):

- (i)  $(\forall x) (\exists y) (\text{boy } x \rightarrow (\text{girl } y \wedge \text{loves } (x, y)))$
- (ii)  $(\exists x) (\forall y) (\text{girl } x \wedge (\text{boy } y \rightarrow \text{loves } (y, x)))$


SD1<sub>(2c)</sub> [S [DP Every boy]<sub>i</sub> [DP some girl]<sub>j</sub> [<sub>t<sub>i</sub></sub> loves <sub>t<sub>j</sub></sub>]]

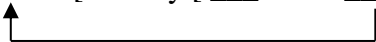
SD2<sub>(2c)</sub> [S [DP Some girl]<sub>j</sub> [DP every boy]<sub>i</sub> [<sub>t<sub>i</sub></sub> loves <sub>t<sub>j</sub></sub>]]

The structural descriptions given here serve to distinguish the different scope relations that can hold between the two quantifier phrases, but they both deviate significantly from the surface form of (2c). In each case, the DPs that occupy the external and internal argument positions of ‘love’ both undergo movement (via *quantifier raising*) to higher nodes, leaving behind a pair of *trace* elements, *t<sub>i</sub>* and *t<sub>j</sub>*, that are marked by indices that serve to link the raised DPs with their initial argument positions. Depending on the order in which quantifier raising occurs, we thus get two different structural descriptions for (2c).

Both of the cases of ambiguity considered thus far provide striking examples of how semantic theories can and must be responsive to features of the underlying syntactic arrangements of sentences, which are not at all apparent at surface form.<sup>27</sup>

Finally, we turn to (2d). On one interpretation of (2d), the duck is ready to eat something, whereas on the other, the duck is ready to be eaten by someone. According to the first interpretation, the DP ‘the duck’ seems to be somehow linked with the subject position of the infinitive verb ‘to eat’, whereas on the second, the same DP is linked with the lower verb’s object position:

SD1<sub>(2d)</sub> [The duck [is ready [ \_\_\_ to eat \_\_\_ ]]]  


SD2<sub>(2d)</sub> [The duck [is ready [ \_\_\_ to eat \_\_\_ ]]]  


The example raises interesting questions regarding the kinds of semantically relevant relations that will be encoded in the LF representations of sentences. For present

---

<sup>27</sup> For a detailed account of how a truth-conditional semantic theory would go about assigning semantic values to structural descriptions containing raised quantifiers and their corresponding trace elements, I refer the reader to Heim and Kratzer 1998.

purposes, given the obvious ambiguity of (2d), we shall take it that our LF representations for (2d) will represent the kinds of structural links that obtain between elements that occupy the different positions in syntactic structure.

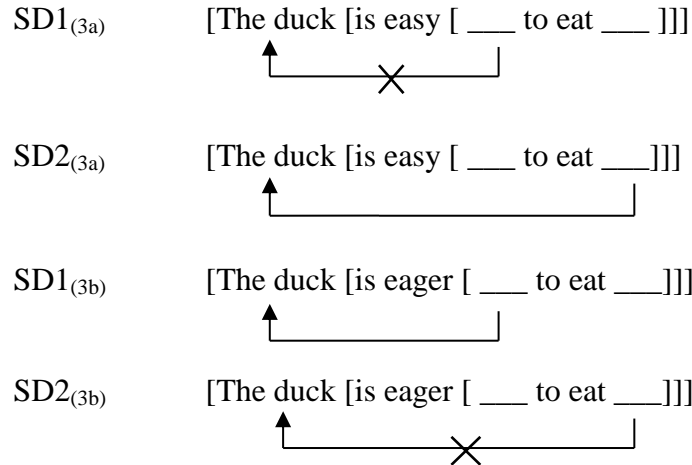
### 3.2. Negative facts

Whilst instances of ambiguity are interesting in the sense that they show that more than one set of truth-conditions can be associated with the same sentence (where we are using the term ‘sentence’ somewhat loosely here to pick out a particular surface form), some of the examples discussed also raise important questions about why the range of interpretations we are able to assign to sentences are sometimes limited in certain non-obvious ways. For example, in relation to (2b), one might well wonder why it is that speakers of English are *not* able to understand the sentence as having a possible interpretation according to which Bill called John, and *Bill* is from London. The fact that speakers are not able to understand the sentence in this way appears to indicate that although the PP ‘from London’ can either be understood as combining with the VP or with its complement NP, there seems to be a constraint that bars it from combining with, and so modifying, the NP subject of the sentence. Such *negative facts* are semantically relevant, as they serve to show that the underlying syntactic relations that hold between the constituent expressions of sentences play a role in determining not only which interpretations speakers *are* able to associate with those sentences, but also which interpretations *are not* available.

Similar observations arise in connection with the sentences of (3):

- (3) a. The duck is easy to eat
- b. The duck is eager to eat

As we saw, if we were to substitute ‘ready’ for ‘easy’ or ‘eager’ in these sentences, the resulting sentence would be ambiguous. As things stand, though, neither (3a) nor (3b) is ambiguous. That is, (3a) can *only* mean that it is easy for someone to eat the duck, and (3b) can *only* mean that the duck is eager for *it*, the duck, to eat something. Hence, we get the following limitations on the range of structural descriptions available for these sentences:



Such cases show how different lexical items, e.g. ‘ready’, ‘easy’, and ‘eager’, impose different kinds of idiosyncratic constraints on the ways in which other expressions they combine with in sentences can be understood to relate to one another. Again, such interactions between the lexical and syntactic features of sentences are something we should expect a complete semantic theory to be able to account for.

Finally, to mention one last case that has not yet been discussed, consider (4):

(4) The men told the women to vote for each other

This sentence means, and so is true iff, the men told each of the women to vote for another woman. This indicates that the determiner phrase ‘the women’ stands in a certain syntactic relation (an antecedent binding relation) with the reciprocal expression ‘each other’. The question to ask though is why this relation can only obtain between these particular parts of the sentence and no others. That is, when we come to explaining how speakers interpret the sentence, we want to be able to say why it is that the sentence *cannot* mean that the men told each of the women to vote for another man, or that the men told the women that they (the men) would vote for another man (Higginbotham 1985, p.549). Once again, this example serves to emphasise the semantic relevance of negative facts.

To sum up, then, I take it that there are good reasons for endorsing Higginbotham’s claim that ‘the data of semantics should be seen in both their positive and their negative aspects: for any given expression, that it *does* mean *X*, or *can* mean *X*, and that it *does*

*not* mean *Y*, or *cannot* mean *Y*, are facts to be deduced in semantic theory' (1985, p.548).<sup>28</sup>

### 3.3. Semantic entailment

One other range of semantic facts we need to consider are those pertaining to certain kinds of truth-involving connections that speakers recognise as holding between sentences. To start with a simple example, all speakers of English know that if the sentence 'Bill is a bachelor' is true, then so too is the sentence 'Bill is an unmarried man'. Knowledge of which expressions are synonymous with each other in this way forms a somewhat obvious, but nevertheless important, part of linguistic understanding. Also, we know that if (5a) is true, then so is each of the following sentences (not to mention others):

- (5) a. Bill danced and Mary laughed loudly
- b. Bill danced
- c. Mary laughed loudly
- d. Someone danced
- e. Someone laughed loudly
- f. Bill did something
- g. Something loud happened

Truth-conditional semantic approaches have proven highly successful when it comes to explaining such entailments. That is, by giving the meanings of sentences in terms of their formally determined truth-conditions, such approaches have been able to apply standard formal analyses to the meanings of sentences that serve to bring out the relevant connections. For instance, (5b) and (5c) both follow from (5a), simply by way of conjunction elimination. (5d) and (5e) are then entailed by existential introduction,

---

<sup>28</sup> For detailed discussion on the kinds of syntactic relations that are involved in the kinds of cases discussed here, see Radford (1988), Haegeman (1994), and Hornstein *et al.* (2005).

which involves generalising over the positions occupied by ‘Bill’ and ‘Mary’ in the logical forms of the previous sentences.

Things are slightly more complicated in the case of (5f)-(5h). To account for such entailments, Davidson proposed that the kinds of expressions that refer to activities and events should be ‘construed as containing a place, for singular terms or variables, that they do not appear to’ (Davidson 1967, p.47). Such variables range over activities and events, and (in most cases) remain implicit, undergoing existential closure in logical form. Modifying Davidson’s original idea, we thus end up with the following representation for the logical form of (5a):

$$(6) \quad (\exists e) [\text{AGENT}(\text{Bill}, e) \wedge \text{DANCING}(e)] \wedge (\exists e') [\text{AGENT}(\text{Mary}, e') \wedge \text{LAUGHING}(e') \wedge \text{LOUD}(e')]$$

Such representations of the semantic logical forms of sentences introduce an array of new concepts that have not yet been discussed. Perhaps most strikingly, we now come to understand the arguments of verbs such as ‘dance’ and ‘laugh’ in terms of the kinds of participation relations that obtain between the activities those verbs denote and the individuals referred to by their arguments; e.g. in (6), Bill is explicitly marked as the AGENT of the dancing activity. Such facts can be taken to be linguistically encoded. That is, linguists have long argued that verbs assign  $\theta$ -roles (e.g. AGENT, THEME, GOAL) to their arguments (see Haegeman 1994, chapter 3 for an introduction to  $\theta$ -roles). More fully, the external argument of a verbal head is assigned the  $\theta$ -role AGENT, whereas its internal argument is typically assigned the  $\theta$ -role THEME. Hence, there is good reason to think that the kinds of relations that are encoded in (6) are linguistically determined, and so will be part of the information encoded in an LF representation for (5a) (for further discussion of these matters, which we shall return to later, see Parsons 1990, Pietroski 2005).

The important point for now is that, in taking (6) to be an accurate representation of the truth-conditional semantic content of (5a), it then becomes possible to explain the entailment relations that hold between that sentence and (5f)-(5h). In the case of (5f), having already isolated the left most conjunct of (5a) to get (5b), we then eliminate the

second conjunct in the logical form of (5b) to get (5f), as indicated in (7) (see Lepore and Ludwig 2007, p.179):

- (7) a.  $(\exists e) [\text{AGENT}(\text{Bill}, e) \wedge \text{DANCING}(e)]$   
b.  $(\exists e) [\text{AGENT}(\text{Bill}, e)]$

(7b) says that there is some event such that Bill was the agent of that event, which is as good a paraphrase as any for ‘Bill did something’. A similar approach applies in the case of (5g), as indicated in (8):

- (8) a.  $(\exists e) [\text{AGENT}(\text{Mary}, e) \wedge \text{LAUGHING}(e) \wedge \text{LOUD}(e)]$   
b.  $(\exists e) [\text{LOUD}(e)]$

A further range of connections that obtain between sentences are those that involve lexical entailment. For instance, consider the following sentences:

- (9) a. Bill broke the vase  
b. The vase broke  
c. Bill hit the wall  
d. \*The wall hit  
e. The men loaded the truck with hay  
f. The hay was loaded on the truck

(9a) entails (9b), but (9c) does not entail the anomalous (9d). (9e) entails (9f), but not vice versa, as (9e) implies that the truck was loaded with hay *till full*, which would not necessarily be the case if the men loaded just one bale of hay onto the truck<sup>29</sup>. These kinds of semantic facts all seem to have something to do with the way speakers understand the meanings of particular lexical items, in that they reveal interesting features of how those lexical items, e.g. ‘break’ and ‘load’, combine with their arguments (see Fillmore 1968; Levin and Rappaport Hovav 1995 for discussion).

---

<sup>29</sup> Whether the verb ‘loaded’ necessarily implies ‘till full’ is disputed, but we will presume here this for the purposes of exposition.



### 3.4. Context sensitivity

The kinds of semantic facts we have looked at thus far all have something to do with ways in which the purely formal, or lexico-syntactic, features of sentences determine and constrain their truth-conditional content. There is a range of semantic facts, however, which pertains to various ways in which extra-linguistic factors can contribute to linguistic understanding. That is, when it comes to determining the truth-conditional content of sentences, speakers are often required to take into account certain pieces of information about features of the contexts in which those sentences are used. This is known as the phenomenon of *context sensitivity*, and it is a phenomenon that has a major bearing on attempts to explain how speakers understand language.

The most obvious, most uncontroversial, way in which context contributes to the determination of content is by fixing the semantic values of standard indexical expressions, where the term *standard indexical expression* refers to expressions such as ‘I’, ‘now’, ‘here’, ‘yesterday’, ‘that’, ‘he’, and so on:<sup>30</sup>

- (10) a. I sang
- b. You laughed
- c. They left

When speakers of English encounter utterances of the above sentence, they know that in order to be able to evaluate those utterances for truth, they are going to have to take into account certain features pertaining to the context of utterance. Consequently, attempts to specify the truth-conditions of such sentences via the use of T-sentence statements of the following form – ‘I sang’ is true iff I sang – will tend to go awry, as there is never any

---

<sup>30</sup> I shall take the set of standard indexical expressions to be roughly that list of expressions set out by Kaplan, which includes ‘the pronouns ‘I’, ‘my’, ‘you’, ‘he’, ‘his’, ‘she’, ‘it’, the demonstrative pronouns ‘that’, ‘this’, the adverbs ‘here’, ‘now’, ‘tomorrow’, ‘yesterday’, the adjectives ‘actual’, ‘present’, and others’ (Kaplan 1977, p.489). I shall only tend to differentiate the demonstratives from this set of expressions when the discussion demands it.

guarantee that the individual referred to by the indexical expression in the sentence mentioned will be the same thing referred to by the individual using the statement to interpret the sentence.

Standard indexicals, then, are expressions whose semantic values are fixed by speakers relative to contexts of utterance. That said, there is an important distinction to be made regarding *how* speakers go about fixing the semantic values of different kinds of indexical expressions. The *automatic* (or *pure*) *indexicals* – e.g. ‘I’, ‘today’, ‘yesterday’ – have their semantic values fixed relative to particular features of context independent of speakers’ intentions. For instance, on uttering ‘I’ at a context *c*, a speaker cannot help but refer to herself, the person performing the utterance at *c*. Similarly, a speaker who believes it to be Monday might utter ‘yesterday’ on a Tuesday, intending to refer to Sunday; but whatever the speaker’s intentions, such an utterance would refer to the day before the day of utterance; i.e. Monday (Perry 1977, p.6).

In contrast, the *intentional indexicals* (which Kaplan refers to as the *true demonstratives*) – e.g. ‘that’, ‘he’, ‘this’ – have their semantic values fixed partly in relation to the referential intentions of speakers who use them (Kaplan 1989, p.582). So, when a speaker A utters ‘He is mad’, although there is clearly *some* constraint on the type of entity A could be referring to with ‘he’, in that the referent must be masculine in some sense, fixing exactly who or what A refers to by using that expression will involve figuring out A’s referential intentions (where an accompanying act of demonstration might provide some clue as to those intentions). If one is unable to identify A’s referential intentions on the relevant occasion, then one’s interpretation of the truth-conditional content of the utterance will remain incomplete (Kaplan 1977, pp.490-491).

It should be noted that the distinction between automatic and intentional indexicals is not always clear-cut. For instance, ‘now’ has typically been taken to refer automatically to the time of its utterance, and ‘here’ to the location of its utterance. However, the duration of time that a speaker intends to refer to with an utterance of ‘now’, or how large an area a speaker intends to refer to with an utterance of ‘here’, might vary on different uses of such expressions (Perry 1998, p.8). To illustrate, one might utter the sentence ‘I like it here’ intending to express one’s liking for the room one is in, or the

building, or the city, etc. This seems to suggest that ‘now’ and ‘here’ are part automatic, in that the duration or area referred to by uttering either expression must contain the instant in time or the location at which the utterance takes place, but are also part intentional, in that the exact duration or area referred to is identified relative to speakers’ intentions.

To further complicate things, though, utterances of both expressions can be used to refer to times and locations other than those at which their uses take place. For example, talking about a long ago battle in a faraway place, one might say ‘The rebels occupying the town knew that they must make their stand right here right now’, and in so doing refer to the very place and time at which the rebels made their stand. What this example seems to indicate is that speakers can *shift* perspective from an *utterance context* to something like a *narrative context* when it comes to fixing the values of *some* indexical expressions, interpreting occurrences of such expressions as if from the perspective of some narrative context (see Predelli 1998). Acknowledging this need not involve giving up on the idea that speakers fix values for, say, ‘here’ and ‘now’ relative to locations and times at contexts, but in showing that there can be some question as to *which* context is relevant when doing so, we are led one step further away from the idea that such these expressions automatically refer to the times and locations of the contexts of utterance at which they are used.

Let us now consider how a semantic theory might account for this. Setting aside certain complications, one thing that seems to be clear is that each occurrence of an indexical expression is assigned its semantic value relative to some specific feature of context. On the basis of this observation, Davidson proposed that the T-sentences generated by a semantic theory would specify the truth-conditions of sentences relative to a fixed set of contextual parameters; i.e. *speakers* and *times*. The axioms for indexical expressions would state that determining the semantic values of indexical expressions would involve identifying speakers and times relative to contexts of utterance (Davidson, 1973, p.131)

The insight was developed in much more detail in the work of Kaplan (1977), who came up with a way of providing context-independent specifications of the meanings of all indexicals. On Kaplan’s theory of indexicals, the meaning of any expression is specified

in terms of its context-invariant *character*, which is represented as a function from contexts to contents. A context, on this view, is characterised in a purely formal manner as an ordered set of contextual parameters that minimally include those ‘brute’ features that make up any context of utterance, i.e. *speaker*, *time*, *location*. Content is characterised as a function from circumstances of evaluation to extensions; e.g. the content of a sentence is a (structured) proposition, which, relative to a circumstance of evaluation, takes a truth-value as its extension.

Kaplan’s claim was that whilst most expressions have *fixed* characters, in that they take the same content as value at all contexts, the characters of indexical expressions are *context sensitive*, in that the contents they take as values vary with the values of the contextual parameters with which they are associated. For example, the character of ‘I’ can be represented as a function from any context of utterance *c* to the value of the speaker parameter at *c*, which is just to say that an utterance of ‘I’ at a context *c* has as its content the speaker who performed the utterance at *c*. As for the intentional indexicals, for present purposes we can take it that the semantic values of expressions such as ‘you’ and ‘that’ are also fixed relative to contextual parameters, say for *addressees* and *demonstrata*, where the values of those parameters at particular contexts of utterance are identified on the basis of accompanying *acts of demonstration* (Kaplan 1977), or relative to the *directing intentions* of speakers (Kaplan 1989).<sup>31</sup>

The advantage of Kaplan’s theory is that it shows how the context-independent meanings of indexical expressions (i.e. their context-invariant characters) determine their contents at any context. By incorporating the notion of character into a truth-

---

<sup>31</sup> Following Kaplan, one might take it that there are contextual parameters for *sequences* of addressees or demonstrata, so as to deal with cases where multiple intentional indexicals occur in a sentence; e.g. ‘That is smaller than that’, or ‘I’m not talking to you, you, or you’. In such cases, the intentional indexical is associated with a syntactically generated index that corresponds to a position in the sequence. For critical discussion of Kaplan’s account of the characters of the intentional indexicals, which he refers to as *true demonstratives*, see Braun 1996, Salmon 2002.

conditional semantic theory, we arrive at a way of specifying truth-conditions for *sentences relative to contexts* via the use of *conditionalized* T-sentences (Davidson 1973; Higginbotham 1988). For example, the truth-conditions for (11) at any context are given by  $CTS_{(11)}$ :<sup>32</sup>

(11) I bought this

$CTS_{(11)}$  If  $u$  is an utterance of ‘I bought this’, and the speaker  $x$  of  $u$  refers with  $x$ ’s utterance of ‘this’ to  $y$  and nothing else, then  $u$  is true iff  $x$  bought  $y$ .

Here, the antecedent of the conditional states conditions for identifying the objects to which the indexical expressions ‘I’ and ‘this’ will refer at any context of utterance in terms of the characters of those expressions, whereas the consequent gives a standard specification of (11)’s truth-conditions for objects thus identified relative to contexts. Crucially, whilst information pertaining to the characters of the indexical expressions contained in (11) does not appear anywhere in the specification of the sentence’s truth-conditions, it is nevertheless represented within the conditionalized T-sentence as part of what speakers know about the meanings of such expressions; i.e. as that bit of information that enables them to secure the referents of the indexicals at contexts. As a result, we have a way of specifying the truth-conditions of (11) at any context of utterance without having to spell out of how speakers actually go about identifying the referents of the indexical constituents of (11) relative to  $c$ .

A semantic theory represents a context of utterance in a purely formal, abstract manner, as an ordered set of objective contextual parameters. The truth-conditions of sentences are then specified relative to contexts of utterance so represented. The context-invariant meaning (or character) of each indexical expression encodes the contextual parameters relative to which values of its token occurrences are fixed. The theory makes the relevant information explicit by specifying the truth-conditions of sentences that contain indexicals in terms of conditionalized T-sentences. Such T-sentences specify the truth-conditions of the relevant sentences independently of how the semantic values of the

---

<sup>32</sup> In chapter 3, we shall look at how Borg (2004) develops this approach.

indexical expressions they contain are secured (Higginbotham 2006). In what follows, we will have cause to raise further questions about standard indexicals, but generally speaking, I shall take it that by adopting the account of standard indexicality sketched here, a semantic theory would explain how the meanings of such expressions determine the contributions they make to the truth-conditions expressed by utterances of sentences in which they occur.

#### 4. Summary

In this chapter my aim was to explore a particular conception of what a semantic theory for a natural language would look like. To that end, I explored the theoretical foundations for truth-conditional semantics as they emerged in the work of Davidson, I showed how we might go about the task of constructing a semantic theory that is true to the truth-theoretic framework Davidson developed, and I assessed how such a theory fares when it comes to the task of accounting for a variety of semantic facts. As a result, I hope to have shown that a truth-conditional semantic theory of the kind envisaged by Davidson offers a rich and systematic method of interpretation for a natural language. Indeed, I hope to have shown why the truth-conditional approach has become the standard approach in contemporary semantics.

That said, though, the rest of the thesis will explore a problem truth-conditional semantic theories face; *viz.* the problem of the phenomenon of pervasive context sensitivity. We have seen that a truth-conditional semantics can accommodate instances of context sensitivity that arise from occurrences of standard indexical expressions, but in the next chapter, we shall look at other ways in which features of extra-linguistic context impact on the determination of truth-conditional content. The question to ask is whether those further varieties of context sensitivity can be accommodated by a truth-conditional semantic theory.

## CHAPTER 2

### PERVASIVE CONTEXT SENSITIVITY

In the previous chapter, we looked at a variety of facts regarding the ways speakers interpret the truth-conditional semantic content of natural language sentences, and we showed how the truth-conditional semanticist accommodates such facts so as to defend the claim that the formal, lexico-syntactic features of sentences, fully articulated at LF, determine complete sets of truth-conditions.

As we saw, though, there are some cases where certain *extra-linguistic* features of context must be taken into account in order to determine a sentence's truth-conditions. In relation to such cases, which arise from occurrences of standard indexical expressions, it was shown that we can account for the role context plays in determining content by arguing that the meanings of such expressions encode information regarding the kinds of contextual factors speakers must take into account in order to determine what contributions they make to content. So, although the phenomenon of context sensitivity indicates that speakers must look beyond the purely formal features of sentences in order to determine their content, and therefore that we determine the content of sentences *relative to contexts*, it seems that a semantic theory can nevertheless account for this without having to abandon the assumption that content is linguistically determined.

However, this conclusion may turn out to be premature, and in the following section I look in more detail at the phenomenon of context sensitivity, and consider whether there are other ways in which extra-linguistic features of context contribute to the determination of content.

#### 1. Beyond standard indexicality

I now want to show that there are good reasons for thinking that the phenomenon of context sensitivity is not restricted to instances of standard indexicality. This outlook has been adopted by many contemporary philosophers of language, and there have been

various arguments set forth in its defence. I shall refer to such arguments as *context shifting arguments* (CSAs), and in this section I shall show how CSAs have been used to provide support for the claim context plays a more pervasive role in determining content than that of fixing values for standard indexicals.

### 1.1. Context shifting arguments

CSAs are designed to show that ‘sentences are compatible with different conditions for their truth’ (Travis 2008, p.112); i.e. are apt to express different sets of truth-conditions when used at different contexts of utterance. Constructing such an argument for a sentence involves devising a pair of contexts in which utterances of the sentence would differ in truth-value despite being evaluated for truth relative to the same circumstance of evaluation. For example, consider (1):

(1) Bill is ready

Picture Bill in the library preparing for a test. Under such circumstances, we can imagine that an utterance of (1) would count as true in the context of a discussion regarding Bill’s chances of passing the test, but false in the context of a discussion about whether Bill is dressed for an evening out. This indicates that the conditions that must obtain in order for the sentence to count as true are apt to vary across contexts of utterance. Hence, although ‘ready’ is not a standard indexical, it appears that it is context sensitive, in that the contribution it makes to the content expressed by utterances of (1) varies across contexts. The same kind of reasoning can be applied to each of the following sentences:

- (2) a. Every student smokes
- b. Bill is tall
- c. Mary’s horse won the steeplechase
- d. It is raining

The truth-values speakers assign to utterances of these sentences may vary even if those utterances are evaluated for truth relative to the same circumstances of evaluation. Hence, speakers will assign different truth-values to utterances of (2a) depending on



which *quantifier domain* is contextually salient ('every student *in Bill's class*'/ 'every student *in the UK*'); to utterances of (2b) depending on which *comparison class* is contextually salient ('tall *for a twelve year old*'/'tall *for a basketball player*'); to utterances of (2c) depending on which *possession relation* is contextually salient ('the horse *Mary owns*'/ 'the horse *Mary bet on*'); and to utterances of (2d) depending on which *location* is contextually salient ('raining *in Paris*'/ 'raining *in Stockholm*').

Such cases are simple, but they serve to illustrate that there are a variety of expression types, e.g. quantifier phrases, comparative adjectives, possessives, weather verbs, etc., whose contributions to content can only be properly determined once certain specific pieces of contextual information have been taken into account; e.g. information about contextually salient domains, comparison classes, possession relations, locations, and so on.

In other putative cases of context sensitivity, it is harder to say with any degree of generality precisely what kind of contextual information speakers must look to in order to understand the content of the sentence uttered. For instance, Travis (e.g. 1996, 2008) has discussed a variety of cases where it is certain peculiar features of particular contexts of utterance that bring about the possibility of variation in a sentence's truth-conditional content:

- (3) a. The leaves are green
- b. The milk is in the fridge
- c. The ink is blue

Focusing on (3a) for the time being, Travis presents us with the following CSA:

A story. Pia's Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, 'That's better. The leaves are green now.' She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. 'The leaves (on my tree) are green,' Pia says. 'You can have those.' But now Pia speaks falsehood. (Travis 2008, p. 111)

What Travis's examples are meant to illustrate is that the truth-conditional contents of sentences vary not only relative to what might be thought of as generally specifiable contextual parameters (e.g. domains, locations, comparison classes, etc.), but that the diverse practices and interests that engage and motivate speakers when they put sentences to use on different occasions can also be shown to have an effect on how we understand the truth-conditional contents expressed by utterances of sentences.

For example, if your purpose is to find some milk for your cereal, then you will take a speaker who utters (3b) to say something false if there is only a puddle of milk in the fridge. That said, if your purpose is to find out whether the fridge needs cleaning, one might then take a speaker to say something true by uttering that sentence in such a set of circumstances. Similarly, you might take it that ink that looks black in a bottle counts as blue ink if it appears blue when scribbled on paper, but then you might take it not to count as blue ink if you are entirely interested in how it looks whilst in the bottle.

There are various questions to be asked about the extent to which CSAs can be convincingly applied. Some have taken CSAs to be convincing only when applied to a very narrow range of sentences, whereas others (e.g. Travis) have argued that there is no obvious limit on the range of sentences to which CSAs can be applied, and what is more, no obvious limits on what kinds of contextual factors might be shown to make a difference when it comes to determining the truth-conditions a given sentence might be used to express. For now, though, I am merely interested in the possibility that there may be certain instances of context sensitivity that are not triggered by occurrences of standard indexical expressions.

Ultimately, I take it that the above examples reveal an interesting range of facts about the ways speakers interpret the truth-conditional contents of the sentences discussed thus far. That said, there is an important distinction to be made regarding the way these examples operate. That is, in the case of 'Bill is ready', it appears that context serves to provide us with information that is not explicitly specified in the original sentence, but which is nevertheless essential for determining whether the sentence is true relative to any given context of utterance; i.e. it seems to be necessary to find out what activity Bill

is ready to perform or engage in, as it is hard to imagine what kinds of conditions would have to obtain in order for it to be true of Bill that he is ready *simpliciter*.

In the case of ‘Every student smokes’, though, the appeal to contextual information appears to serve the role of making the truth-conditional content of the sentence more specific; i.e. at face value, the sentence seems to state that for any individual anywhere, if that individual is a student, then he or she smokes, but we take it that when a speaker utters the sentence, they intend to say something about a particular domain of students.

Following Bach (1994), I shall refer to the former kinds of cases as cases of *semantic incompleteness*, and the latter as cases of *semantic expansion*. I now want to consider why the idea that context might make these different kinds of contributions to content raises issues for standard truth-conditional semantic theories.

## 1.2. Semantic incompleteness

The idea that the compositionally determined semantic content of a sentence does not always determine a complete set of truth-conditions seems to pose a significant challenge for truth-conditional semantic theories. From the outset, we have taken it that the meanings of sentences can be specified in terms of the conditions that must obtain for those sentences to be true, and so that we can analyse the meanings of simple expressions by looking at the way they contribute to the truth-conditions of the sentences in which they occur. However, although we know what the constituent expressions of the sentences listed in (4) mean, and we take it that we can provide accurate descriptions of the structural relations that obtain between the expressions which those sentences comprise, it would appear that knowing this much is not enough to put us in a position to be able to judge whether utterances of the sentences would be true or false relative to any given circumstances of evaluation:

- (4) a. Bill is ready
- b. Bill is tall
- c. It is raining

Putting the point somewhat differently, there are certain pieces of information that must be taken into account in order to determine the truth-conditions of these sentences that are not explicitly specified by any of their constituent expressions. For example, there is no constituent of (4c) that takes a location as its semantic value. Hence, it would appear that, in and of themselves, the lexico-syntactic features of the sentences of (4) provide only partial, or incomplete, representations of the truth-conditional content speakers associate with utterances of those sentences. In order to arrive at a complete set of truth-conditions, then, a speaker must utilise her knowledge of context in order to identify the missing pieces of information, and so determine a complete set of truth-conditions.

The problem that the phenomenon of incompleteness raises for a truth-conditional semantic theory is this: if context can play a constitutive role in determining what truth-conditional content an utterance of a sentence expresses, then there is no reason to think that linguistic meaning alone determines those semantic relations that speakers recognise as holding between expressions and the world. That is, once we take it that the relations that obtain between expressions and the world are, at least to some extent, contextually determined, then it seems we must accept that linguistic meaning in and of itself does not determine stable, shared relations between sentences and the world. Instead, those relations are, at least in part, forged on the basis of some form of pragmatic competence.

The worry that this gives rise to, though, is that there seem to be no obvious constraints on what kinds of pragmatic factors, what kinds of features of contexts, speakers might take to be relevant when it comes to determining the content of an utterance of a given expression. It is hard to imagine how one would go about the task of providing a systematic account of the way speakers use language to talk about the world that is based on an appeal to pragmatic competence.

The truth-conditional semantic theorist, motivated by such concerns, will thus attempt to provide an account of putative instances of incompleteness that ultimately shows that, despite appearances, contextual factors do not play a constitutive role in the determination of content. Prominent accounts that aim to accomplish precisely this will be analysed in detail in subsequent chapters.

### 1.3. Semantic expansion

Semantic expansion occurs when speakers take the compositionally determined truth-conditional content of a sentence to differ from the content intuitively expressed by an utterance of that sentence. In such cases, the intuitive truth-conditional content of the sentence is an ‘enriched or elaborated version of the [content] explicitly expressed by the utterance itself’ (Bach 1994, p.133). Some form of expansion is taken to occur in each of the following cases:

- (5) a. Every student smokes
- b. You’re not going to die
- c. Bill has not eaten breakfast

In each of these cases, it seems clear that a complete truth-conditional content is determined by the formal features of each sentence. However, when speakers utter such sentences on particular occasions of use, speakers will typically take them to have expressed something more specific than the sentence’s compositionally determined content. We have already dealt with how this arises in connection with (5a), and utterances of (5b) will typically be interpreted as being true just if the addressee is not going to die *from some contextually salient cause* (e.g. ‘You’re not going to die *from that cut*’). Similarly, utterances of (5c) will typically be taken to be true iff Bill has not eaten breakfast *today*. It should be noted that these kinds of interpretive effects are not ‘one off’ effects brought about at peculiar contexts of utterance, but reflect general facts about the way speakers actually interpret utterances of the relevant expressions.

Now, there is a sense in which instances of expansion are less of a challenge for the standard semantic theorist than instances of incompleteness. In cases involving incompleteness, there is a sense in which context plays an indispensable, constitutive role in the determination of content, whereas in cases of expansion, the kind of contribution context makes is merely optional. So, the truth-conditional semanticist might simply choose to treat expansion as some kind of optional, *post-semantic* process that does not contribute to the determination of an utterance’s compositionally determined content.

That said, though, instances of expansion give rise to important questions regarding what it is truth-conditional semantic theories actually aim to account for. If the truth-conditions speakers associate with utterances of sentences are sometimes sets of truth-conditions arrived at through some form of expansion, then we have to ask whether a semantic theory (which is, after all, a systematic method of interpretation for sentences in a language) should account for instances of expansion. The problem is that if we *do* take it that a semantic theory should account for the content speakers intuitively associate with utterances of a sentence, then we are once again led to the view that context plays a constitutive role in the determination of content, which in turn raises the potentially significant challenges for lexico-semantic theorist noted earlier.

## 2. Three responses

It is evident that CSAs raise a significant challenge for the view that the lexico-syntactic features of sentences, assessed relative to contexts of utterance, always serve to determine determinate sets of truth-conditions. That said, there are two questions that need to be asked so as to gauge the kind of response that CSAs merit. The first is whether intuitions elicited by such arguments *should* be accounted for by semantic theories for natural languages, i.e. whether they are semantically relevant. The second is whether such intuitions *can* be accounted for within truth-conditional frameworks for constructing such theories. The remainder of the section focuses on distinguishing between the following three responses to these questions: (i) the intuitions elicited *are* semantically relevant, and *can* be accommodated by truth-conditional semantic theories; (ii) the intuitions elicited *are* semantically relevant, but *cannot* be accommodated by such theories; and (iii) the intuitions elicited *are not* semantically relevant, so the fact that such intuitions cannot be accommodated by truth-conditional semantic theories is immaterial. Adopting the terminology of C&L – using the term *contextualism* for any approach according to which the phenomena highlighted by arguments for pervasive context sensitivity are semantically relevant – we can refer to those who endorse,

respectively, (i) as *moderate contextualists*, (ii) as *radical contextualists*, and (iii) as *semantic minimalists*.<sup>33</sup> Let us now look in turn at each of these responses to CSAs.

### 2.1. The moderate response

Moderate contextualists (e.g. Perry 1986, Bach 1994, Szabo 2001, Stanley 2007) take CSAs to be convincing only when applied to a narrow range of sentence types. On this view, all genuine instances of context sensitivity are triggered by a select set of sub-sentential expression types (e.g. quantifier phrases, comparative adjectives, weather verbs, etc.) whose occurrences give rise to the apposite effects on interpretation. Accordingly, it is argued that to account for those instances of context sensitivity that go beyond standard indexicality, we need to inquire into the intrinsic semantic properties of the salient expression types. Moreover, moderate contextualists argue that this can be done by implementing methods more or less in keeping with the basic assumptions of truth-conditional semantics. Such methods can be broadly characterised as follows. Start by establishing that a sentence is susceptible to CSAs, being sure to rule out effects brought about by standard indexicality, tense, ambiguity, polysemy, or any merely seemingly semantic effects. Standard examples include the following:

- (6) a. Mary is tall
- b. Every student failed
- c. It is raining

Next, via careful reflection on intuitions elicited by CSAs, identify the expression responsible for the context sensitivity of the host sentence and the contextual factor (or parameter) relative to which the semantic value of the expression varies. For the sentences of (6), this yields the following: in (6a), the semantic value of the comparative

---

<sup>33</sup> The terminology here is adopted from C&L (2005). Borg (2006) opts out of using these terms as they ‘fail to capture the true divide in the contextualist camp’. That said, I shall take it that the dialectic between contextualism and minimalism can be properly set up using C&L’s terminology, just as long as we are careful to mark important difference with the approaches that fall under these rather general headings.

adjective ‘tall’ varies depending on what comparison class is salient at each context of utterance (e.g. tall *for a five year old*); in (6b), the value of the quantifier phrase ‘every student’ varies depending on the domain to which the quantifier is restricted at each context (e.g. every student *in Bill’s class*); and, in (6c), the value of the weather verb ‘rain’ varies depending on the contextually salient location of the raining event (e.g. raining *in Paris*). Finally, having identified the context sensitive expression, posit a variable-like element that is intrinsically associated with that expression, and whose value varies across contexts of utterance relative to the appropriate contextual factor. The resulting view is that the propositions expressed by context sensitive sentences vary systematically relative to the values the putative variable-like elements take across contexts.

Advocates of different versions of moderate contextualism adopt different outlooks regarding how best to characterise the pertinent variable-like elements. Stanley (2000), for instance, argues that such elements are aponic pronominals that are situated in the ‘real’ logical forms of sentences (i.e. are represented at LF), but which have no morphological reflex at surface form (i.e. are not represented at PF). In contrast, Perry (e.g. 1986, 2001) argues that such elements can be characterised as *argument roles*, which are not associated with expressions as such, but rather with the relations they denote. On this view, the contextually determined elements that saturate argument roles are *unarticulated constituents*; i.e. elements of propositional content that are not *articulated* at any level of linguistic representation. As for others, Bach (1994) characterises the salient elements in terms of extra-linguistic *conceptual gaps*; Taylor (2007) in terms of lexically (but not syntactically) represented *subs syntactic parameters*; Rothschild and Segal (2010) in terms of *indices* on lexical items; and so on. These variations on the moderate contextualist outlook are divided, and some are directly opposed to others, but generally speaking, they all adopt the same basic procedure for accounting for non-standard instances of context sensitivity.

Moderate contextualists, then, are committed to the claim that the context-independent meanings of context sensitive expressions can, in principle, be specified in such a way as to fix (more or less precisely) what types of contextual information must be



accounted for to determine the contributions such expressions make to the propositions expressed by utterances of sentences in which they occur. The resulting view is that sentences containing context sensitive expressions do not express complete propositions independently of context, but *propositional radicals, fragments, or schemas*; i.e. semantic representations permeated by specifiable types of gaps or slots that must be saturated relative to contexts to attain complete propositions.

## 2.2. The radical response

Radical contextualists (e.g. Searle 1978, Travis 2008, Recanati 2010) adopt a very different outlook. On their view, there is no obvious limit on the range of sentence types to which CSAs can be applied; so, *if* we accept the force of such arguments, we will ultimately be led to accept that the phenomenon of context sensitivity is not restricted to occurrences of any exclusive set of expression types. As a consequence, radical contextualists reject the idea that moderate contextualist methods are fit for accommodating the full range of intuitions CSAs elicit.

To explain why, consider the following example. Adopting a moderate contextualist stance, let us assume that:

(7) Bill is ready

expresses, say, a propositional fragment that must be saturated by activities Bill is ready to perform or engage in relative to contexts of utterance. If this is so, then rather than uttering (7) at a context *c*, one might instead utter a more elaborate sentence that explicitly specifies the activity salient at *c*, and so express a complete proposition.<sup>34</sup>

(8) Bill is ready to eat something

---

<sup>34</sup> Moderate contextualists tend to be in agreement on the idea that implicit contextually determined elements of propositional content can be made explicit so as to produce sentences that express complete propositions. Coming up with a sentence that expresses a complete proposition, or thought, is more or less a matter of being somewhat prolix. See Bach 2006.

The problem for the moderate contextualist is that even though (8) makes explicit reference to the activity Bill is ready to perform at *c*, it can be shown that the sentence does not determine a complete proposition that is *the* content expressed by all its utterances. To illustrate, picture Bill asleep after an operation. Carrying out a test on Bill as he sleeps, the doctor utters (8), and so says something true; i.e. the test confirms Bill is showing signs of being ready to eat something (low blood-sugar levels, etc.), a finding worthy of comment from the doctor's perspective. The nurse repeats the doctor's utterance to the cook, who promptly takes a meal to Bill's room, but finding Bill asleep, the cook judges the nurse to have said something false; i.e. Bill is not awake yet, and so is *not* ready to eat something. The conditions that must obtain for Bill to count as *being ready to eat something* vary, then, depending on the contextually salient interests that motivate the doctor and the cook.<sup>35</sup>

This raises complications for the view that it is possible to provide context-independent specifications of the types of contextual information required in order to determine the propositions expressed by utterances of (7), as even when we elaborate (7) to produce another sentence, (8), in which the relevant information is made explicit, further

---

<sup>35</sup> Note that the example does not play on the ambiguity between *psychological* and *material* readings of 'ready'. Such an ambiguity is clear in the sentence 'The duck is ready to eat something'. On the psychological reading, the duck is ready to eat in the sense of *wanting* to eat something, whereas on the material reading, it is ready to eat in the sense of having properly been cooked (Landau 1999). In the CSA presented above, I take it that the material reading is operative in both contexts. What is more, I have attempted to construct the CSA in a way that it eludes explanation in terms of some form of stage-level/individual-level ambiguity on 'ready'. The stage-level reading is operative in both contexts (the doctor takes Bill to be in the temporary state of being ready to eat something, a state that will change once he has eaten his fill, the cook takes Bill *not* to be in such a state). The fact that 'ready' admits psychological/material *and* stage-level/individual level ambiguities only goes to show what a complex semantic life the adjective leads.

contextual factors can be shown to play a constitutive role in determining the propositions expressed by utterances of *that* sentence.

One option open to the moderate contextualist to overcome this difficulty is to deny that the CSA elicits intuitions for the context sensitivity of (8). However, for this to work, the moderate contextualist would have to explain why *this* CSA is less convincing than those upon which her own position relies. Alternatively, the moderate contextualist can argue that (8), like (7), harbours a variable-like element whose value varies relative to some specific contextual factor. Such a response, however, faces its own difficulties. First, radical contextualists will argue that for any contextual factor one might specify, ‘further possible factors would yield more than one distinguishable thing to be said for fixed values of [the specified factor]’ (Travis 2008, p.115). The following example further illustrates the point:

A story. Pia’s Japanese maple is full of russet leaves. Believing that green is the colour of leaves, she paints them. Returning, she reports, ‘That’s better. The leaves are green now.’ She speaks truth. A botanist friend then phones, seeking green leaves for a study of green-leaf chemistry. ‘The leaves (on my tree) are green,’ Pia says. ‘You can have those.’ But now Pia speaks falsehood. (Travis 2008, p. 111)

On Travis’s view, here we have two occasions at which different understandings of what counts as being green are in play. The difference can be described straightforwardly: on one occasion, leaves painted green count as green, whereas on the other, only naturally green leaves so count.

A moderate contextualist might posit a variable-like element in the semantics of ‘green’ to account for the difference. Szabo (2001), for instance, argues that ‘green’ is associated with a variable that ranges over parts of objects. On the first occasion described in the Travis case, the value of the variable is fixed in such a way that the predicate applies to the surfaces of the leaves (i.e. the leaves need only be *superficially* green), whereas on the second, it is fixed in such a way that the predicate applies to the whole of each leaf (i.e. they must be green *through and through*). However, Travis

argues that there are indefinitely many ways in which understandings of how leaves must be to count as being green might vary across further occasions of use:

Suppose the leaves were not painted (or were painted red), but had a fluorescent green mould growing on them. Or suppose they are painted, but in pointillist style: from a decent distance they look green, but up close they look mottled. Is that a way of painting leaves green? It might sometimes, but only sometimes, so count. So there would be two distinct things to be said in the presumed 'paint counts' sense of 'is green'. And so on. (Travis 2008, p. 112)

Ultimately, the idea is that if there is no limit on the range of CSAs one can come up with (and the radical contextualist sees no good evidence for imposing such a limit), moderate contextualists will be forced to keep positing ever more variable-like elements to cope with the full range of intuitions such arguments elicit.

The second problem is that it is not at all clear whether it is possible to specify with any degree of generality what type of contextual factor varies between the utterances of the doctor and the cook. Indeed, it is hard to think how one would specify exactly what varies without reiterating those details peculiar to the *sleeping Bill* case.<sup>36</sup> Hence, it seems that the specific features of the contexts at play in the relevant case, rather than any generally specifiable features of contexts *as such*, are what cause the content expressed by utterances of the sentence to vary. So, the moderate contextualist method of fixing some general contextual factor relative to which the semantic value of some expression in (8) varies across *all* contexts is not applicable here.

---

<sup>36</sup> The different understandings of 'ready' at play might be described as follows: the doctor understands 'ready' in the sense of *being in the right physical condition to perform the activity*; the cook understands 'ready' in the more restricted sense of *being actively able to perform the activity*. Does this identify a particular type of contextual factor? Do all uses of 'ready' have to take values relative to such a factor; e.g. does that factor come into play when deciding if Bill is ready to start his book, to jump off a cliff, to meet his maker, for his close up, etc.?

The radical contextualist thus claims that there can be contextual effects on the content expressed by utterances of a sentence that are ‘not triggered by an expression in the sentence but take place for purely pragmatic reasons – in order to make sense of what the speaker is saying’ at those contexts. Such effects are ‘free’ in that ‘they are not mandated by the linguistic material but respond to wholly pragmatic considerations’ (Recanati 2010, p.12). Radical contextualists thus reject the view that the meanings of expressions can be specified in such a way as to fix precisely those ways in which the semantic contributions of those expressions will be apt to vary across contexts. Consequently, since all expressions can in principle be shown to be context sensitive, it follows that meaning does not itself determine what semantic contributions expressions make on occasions of use:

It seems to be fairly generally realized nowadays that, if you just take a bunch of sentences...impeccably formulated in some language or other, there can be no question of sorting them out into those that are true and those that are false; for...the question of truth and falsehood does not turn only on what a sentence is, nor yet on what it means, but on, speaking very broadly, the circumstances in which it is uttered. Sentences are not *as such* either true or false. (Austin 1962, pp.110-111)

...in general the meaning of a sentence only has application (it only, for example, determines a set of truth conditions) against a background of assumptions and practices that are not representable as a part of meaning. (Searle 1980, p. 221)

What words mean plays a role in fixing when they would be true; but not an exhaustive one. Meaning leaves room for variation in truth conditions from one speaking to another. (Travis 1996, p. 451)

Abandoning the idea that contextual effects on content are triggered by the intrinsic properties of expressions, radical contextualists appeal to our broader pragmatic competence in order to explain how we are able to keep track of the propositions expressed by utterances of sentences across contexts. For example, Recanati (2010)

argues that simple and complex expressions alike are apt to undergo some or other form of *modulation* when put to use at particular occasions, where ‘modulation’ refers to a family of pragmatic processes that operate on the context-independent meanings of expressions so as to cause their semantic values to shift. Importantly, modulation is never mandated by the intrinsic semantic properties of expressions, but is freely applied by speakers in order to figure out what is said; i.e. it is not a mandatory *bottom-up* process (like saturation), but an optional *top-down* process (Recanati 2010, p.14). Similarly, Travis argues that the properties of words as such do not determine what contributions speakings of them make to content across contexts. Rather, our *parochial sensibility* to features of occasions of use plays an essential role in arriving at appropriate understandings of which such contributions words make (Travis 2006, pp.129-139).

### 2.3. The minimalist response

Semantic minimalists agree with radical contextualists about the consequences of acknowledging the force of CSAs. Cappellen and Lepore (C&L), for instance, argue that moderate contextualism is ‘not a stable position’, as the kind of evidence moderate contextualists rely on to support their views ‘leads directly to radical contextualism’ (C&L 2005, p.39). Likewise, Borg insists that if we take the kinds of effects on interpretation highlighted by CSAs seriously, then ‘one has no reason to think that any context independent specification of propositional content will be free from the possibility of running a CSA’ (Borg 2012, p.144).

Despite this point of agreement, though, minimalists find the radical contextualist outlook about the kinds of processes involved in determining the contents of utterances implausible. On their view, if it were true that speakers utilise *free pragmatic processes* in order to track what is said by putting sentences to use across diverse contexts, it would be impossible to account for the systematic manner in which speakers are able to learn and use natural languages. Without any constraints (or, more to the point, without any constraints that issue from the linguistic meanings of expressions) on the kinds of contextual factors speakers might deem relevant when it comes to interpreting each other’s utterances, there would be no way of systematically interpreting what

propositions speakers express by putting sentences to use on particular occasions. This makes it hard to understand how it is that speakers are able to communicate with each other without making continual mistakes, continually misunderstanding one another's utterances. C&L propose that there are two ways radical contextualists can choose to respond to this concern.

The first is to argue that we do not manage to communicate as successfully as we might imagine. Radical contextualists could simply 'bite the bullet' and abandon the notion that understanding one another involves grasping the same propositions expressed by utterances. For C&L this amounts to the suggestion there is no way to 'secure successful communication across contexts' (C&L 2005, p.126). Not only is such an outcome unappealing, but C&L think it is inconsistent. One problem is that to agree with contextualists, one would have to assume that we understand all of the propositions that they express as premises for their position. To do this, one would have to assume that these propositions are stable and that their truth-values are constant across all potential contexts in which they might be encountered. Hence: 'to interpret the sentences that express RC, you have to assume that RC is *not* true' (C&L 2005, p.128). Put another way, C&L maintain that for the sentence 'Radical contextualism is true' to be a general, truth-evaluable claim about language, its truth-conditions should be fixed across contexts; but why should instances of the predicate *is true* be any less prone to CSAs and incompleteness arguments than other expressions? C&L find that contextualists provide no adequate response to this.

The second response is to argue that it is possible to devise a theory that specifies the kinds of pragmatic processes speakers utilise in order to engage in linguistic communication (this is the approach Recanati (2010) adopts in developing his truth-conditional pragmatics). The problem is that even if we are able to specify what kinds of pragmatic processes speakers implement in order to communicate, there is still no systematic way of determining which of these processes (if any) speakers actually utilise at given occasions of utterance.<sup>37</sup> So, radical contextualists who adopt this response still

---

<sup>37</sup> As Recanati (2010) himself points out, modulation is not a systematic pragmatic process...

fail to give a systematic account of how speakers are able to pair particular utterances with propositional contents. Ultimately, then, C&L conclude that contextualism militates against the prospects of developing constructive accounts of the ways in which speakers are able to understand utterances of sentences across diverse contexts of utterance.

The outcome of all this is that if radical contextualism were correct, ‘it would be miraculous if people ever succeeded in communicating across diverse context of utterance’ (C&L 2005, p. 123). But, as a matter of fact, people *do* succeed in communicating across contexts. So, rather than accepting the force of CSAs, and endorsing radical contextualism as a consequence, we must reject the *mistaken assumption* (C&L 2005, p.53), common to all forms of contextualism, that intuitions about speech act content, e.g. those elicited by CSAs, are always relevant for establishing what constitutes semantic content. Indeed, minimalists argue that separating semantic content from speech act content ‘is necessary both in order to do semantics properly and in order to do speech act theory properly (C&L 2005, pp.153-154). As Borg argues:

...there are good reasons to think that it is simply a mistake to require a *semantic* theory to be answerable to intuitions about *speech act* content, for semantic content is one kind of thing (a repeatable, codifiable, rule-governed kind of thing) while speech act content is another kind of thing altogether (a potentially unrepeatable, nebulous, context-governed kind of thing). (Borg 2012 p.15)

According to semantic minimalists, then, the only way to go about the task of constructing systematic semantic theories for natural languages is to cleave to the view that the semantic content of any well-formed declarative sentence S is a stable, context-invariant content expressed by all utterances of S. This is referred to as the *minimal semantic content semantically expressed* by S, and it is argued that it is our ability to grasp the minimal semantic contents that utterances of sentences express that underpins our ability to engage in linguistic communication across contexts.



### 3. Summary

The phenomenon of pervasive context sensitivity raises a considerable challenge for the claim that the formal, lexico-syntactic properties of sentences encode enough information to assign complete sets of truth-conditions relative to contexts of utterance. The phenomenon of semantic incompleteness throws into question the idea that semantic theories can specify complete sets of truth-conditions for sentences purely on the basis of their formal features. The phenomenon of semantic expansion leads us to question whether the specifications of truth-conditions a truth-conditional semantic theory yields puts us in the position of accurately interpreting the sentences speakers use at particular contexts of utterance.

We have seen that there are three different responses we might adopt in light of the evidence for pervasive context sensitivity that emerges from context shifting arguments. The difference between these responses turns on the question of how we should respond to arguments aimed at adducing evidence for context sensitivity. We have seen that of the three responses considered, only two retain any hope of developing a truth-conditional semantic theory; *viz.* the moderate and the minimalist response.

Consequently, given that our aim here is to assess the claim that information derived from the purely formal features of sentences put us in a position to determine the truth-conditions of those sentences relative to contexts of utterance, we need to find out whether either of these responses actually succeed in defending a standard truth-conditional approach to semantics. Accordingly, the next three chapter will be devoted to performing a careful examination of those responses to the phenomenon of pervasive context sensitivity.

## CHAPTER 3

### SEMANTIC MINIMALISM

We have seen that evidence adduced from context shifting arguments presents a very real problem for truth-conditional semantic theories. That said, we also saw that there have been a number of responses to this challenge. The purpose of this chapter is to elaborate on one such response; *viz.* that provided by *semantic minimalism*. The chapter will begin by giving an overview of semantic minimalism, drawing out the core claims endorsed by minimalists. I then go on to examine the specific, and somewhat distinct, arguments raised in defence the approach by looking at the two versions of semantic minimalism that have arisen to date. I will then finish by raising a number of objections against semantic minimalism, addressing each version of the approach in turn before summarising the strengths and weaknesses of the position.

#### 1. The motivation for minimalism

Two versions of semantic minimalism have been elaborated to date: by Cappelen and Lepore (C&L) (2005) and Borg (2004, 2012). Although these versions of the approach differ in important respects, the chief motivation in both cases is to defend a standard conception of truth-conditional semantics in light of the claim that contextual factors make a more pervasive contribution to the determination of content than that of fixing values for standard indexicals.

Semantic minimalists maintain that the arguments in favour of this claim are predicated on the assumption that in order to give adequate explanations as to what constitutes the truth-conditional semantic content of a sentence *S*, a semantic theory must fully account for the various intuitions speakers have about what constitutes the speech act content expressed by *S* at particular contexts of utterance; i.e. intuitions about *what is said, claimed, asserted*, etc., by utterances of *S*. It follows that if speakers can be shown to have intuitions to the effect that the speech act content expressed by an utterance is in part contextually determined, then this must be accounted for when specifying what constitutes the semantic content of the sentence uttered.

According to semantic minimalists, however, the assumption upon which contextualist arguments rely is mistaken (C&L 2004, pp.53-57; Borg 2012, p.103). On their view, the very idea that a proper conception of semantic content must be fully answerable to the various intuitions speakers have about what constitutes the speech act content of their utterances militates against the prospects of developing systematic semantic theories. Indeed, they claim that any approach aimed at accommodating such intuitions would ultimately have to acknowledge such a pervasive degree of context sensitivity in natural language that it would be impossible to explain how speakers are able to routinely interpret sentences across diverse contexts of utterance.

Minimalists thus argue that in order to construct semantic theories capable of explaining certain basic facts about linguistic understanding, i.e. its systematicity and productivity, it is necessary to adopt a minimal conception of semantic content that abstracts away from many intuitions about speech act content. In effect, their position is that any informative semantic theory must be careful to distinguish semantic content proper – i.e. *minimal semantic content* – from speech act content (Borg 2007, p.351).

### 1.1. The core claims of semantic minimalism

Semantic minimalism can be understood in part as an attack on contextualism: on moderate contextualism because once we fully acknowledge the force of the arguments used to motivate the approach, we will ultimately be led to endorse some form of radical contextualism; and on radical contextualism because ‘it fails to account for how we communicate *across contexts*’ (C&L 2005, p.x). Having shown why minimalists adopt this outlook, let us now take an in-depth look at the core claims of semantic minimalism, endorsed by all advocates of the approach.

First, as has already been mentioned, semantic minimalists hold that semantic content is not speech act content. The semantic content of a sentence *S* is the minimal semantic content expressed by all utterances of *S*, rather than any of the diverse, pragmatically enriched contents speakers intuitively associate with particular uses of *S*. This conflicts with the standard assumption, originating in the work of Grice, that the propositions expressed by utterances of sentences can be specified in terms of *what is said, claimed,*

or *asserted* by such utterances. Minimalists argue that what a speaker intuitively says by uttering a sentence often goes beyond the minimal semantic content expressed by that utterance. So, whilst speakers are no doubt usually interested in figuring out what a speaker says by uttering a sentence, just as they are interested in figuring out the implicatures of such utterances, their ability to do both of these things involves the implementation of extra-linguistic pragmatic processes which are not part of their semantic competence *per se*, and so are beyond the remit of semantic inquiry (Borg 2004, p.128; C&L 2005, p. 204).

Second, the minimal semantic content of a sentence is its compositionally determined content; i.e. it is the content determined by the context-independent meanings of the constituent expressions of the sentence and the context-independent compositional rules that specify the semantic effects of combining those expressions as per the sentence's syntactic arrangement. As Borg puts it, semantic minimalism is committed to the claim that there is an 'exclusively lexico-syntactic route to semantic content' (Borg 2012, p.83). Of course, we saw that moderate contextualists also endorse this claim (though in somewhat different ways). The difference between moderate contextualism and semantic minimalism, however, has to do with the kinds of lexico-syntactic features advocates of the approaches attribute to sentences. For reasons discussed, minimalists are distrustful of the idea that we can, say, attribute 'hidden' lexico-syntactic structure to sentences so as to accommodate intuitions elicited by CSAs. Consequently, minimalists argue that in order to determine the semantic contents of sentences we need only assign values to their surface-level constituents, and apply compositional rules in accordance with their surface structure. That is, minimalists tend to argue that in order to determine the minimal semantic contents of sentences we need only take into considerations those lexico-syntactic features of sentences that are apparent at their surface form.

Third, following from the previous claim, semantic minimalists take the range of genuine contextual effects on semantic content to be severely limited. Indeed, one of the main aims of C&L's semantic minimalism is to present arguments to the effect that all genuine instances of context sensitivity are triggered by occurrences of standard indexicals whose meanings can be specified in terms of non-constant functions from a

limited number of objective contextual parameters to contents. Whilst Borg does not think that defending this claim is the primary aim of semantic minimalism, she does take it to be a consequence of the approach: ‘context can only come to affect semantic content when it is called for by something in the lexico-syntactic form of the sentence and the kinds of lexico-syntactic elements which call for contextual input are themselves limited’ (Borg 2012, p.5). Importantly, on both versions of minimalism, it is argued that in order to resist the implausible consequences of contextualism, one must hold that the only kind of semantically relevant context sensitivity is standard indexicality.

Fourth, modulo sentences that contain standard indexicals, the minimal semantic content expressed by utterances of a sentence *S* is a complete proposition, and is thus characterised as the *the minimal proposition semantically expressed* by all utterances of *S*. This is perhaps the central claim of semantic minimalism. The claim is controversial, since as we saw in the previous section, there are various sentences that contain no standard indexicals, but which appear not to express anything truth-evaluable until features of the contexts in which they are used have been taken into consideration. Minimalists insist, however, that contextual information plays no role in determining the semantic content expressed by utterances of such sentences. So, for example, the sentence ‘Bill is ready’ semantically expresses the same minimal proposition across all occasions of use.

Having laid out the motivation for, and the core claims behind, semantic minimalism, I now want to look at the arguments minimalists present for defending them. My aim is to focus on how each version of minimalism defends the view that sentences express minimal propositions that are often not the propositions we intuitively associate with utterances of sentences, that our ability to grasp the minimal propositions

## 1.2. C&L’s semantic minimalism

C&L, *qua* advocates of semantic minimalism, endorse the claim that speech act content is not semantic content. In adopting this stance, however, they do not intend to deny that there are interesting, often informative connections between intuitions about speech act content and semantic content (C&L 2005, p.57). Rather, they claim that the point is that

one has to be very careful when it comes to identifying which intuitions about speech act content are semantically relevant. On C&L's view, the problem with contextualist approaches is that, by focusing exclusively on intuitions elicited by CSAs, advocates of such approaches fail to take into account a wide range of further intuitions that are of considerable relevance to semantic inquiry (C&L 2005, p.151). This leads contextualists to misconstrue the relationship between speech act content and semantic content, and so endorse 'empirically inadequate' accounts of how speakers actually go about determining the propositions expressed by sentences across different contexts of utterance (C&L 2005, p.87).

A crucial fact about language use that contextualists apparently consistently fail to account for is that speakers can produce accurate disquotational indirect reports of one another's utterances across diverse contexts of utterance. To clarify, what a speaker A said by uttering a sentence S at a context *c* can often (though significantly, not always) be accurately and literally reported by a speaker B at a context *c\** by using a sentence of the form 'A said that S'. This can occur even when the *target-context*, *c*, and the *report-context*, *c\**, are radically different, and even when speaker B is unfamiliar with many of many details peculiar to *c*.

It would seem, then, that we can give purely disquotational specifications of the propositions expressed by utterances of sentences across diverse contexts. What is interesting about this aspect of linguistic communication is that speakers not only seem to have clear intuitions to the effect that indirect reports can succeed as a means of conveying what other speakers expressed by uttering sentences at other occasions of use, but they also have clear intuitions that the range of expression types whose occurrences in sentences block the possibility of producing such reports for utterances of those sentences is severely limited. Given this, C&L argue that speakers' intuitions about indirect reportage are highly relevant when it comes to identifying what effects context has on content.

To draw out the pertinent intuitions, C&L devise three distinct *disquotational tests*. The tests not only demarcate what C&L refer to as the *basic set* of genuinely context sensitive expressions, but they also serve the critical end of establishing that any

sentence that does not partly comprise context sensitive expressions expresses a minimal semantic content that is invariant over contexts of utterance. Let us look at each test in turn.

First, the *inter-contextual disquotational indirect report* (IDIR) test is aimed at showing that ‘an expression is context sensitive only if it typically blocks inter-contextual disquotational indirect reports’ (C&L 2005, p.88). To illustrate, take Mary to have uttered both (1) and (2) at *c*:

(1) Bill is ready

(2) I am ready

Since these sentences were uttered by Mary at *c*, someone at another context *c\** might report on her having done so. Let these reports take the following form:

IDIR<sub>(1)</sub> Mary said Bill is ready

IDIR<sub>(2)</sub> Mary said I am ready

IDIR<sub>(1)</sub> is intuitively true for any target-context/report-context pair one might choose. That is, for any *c* and any *c\**, an utterance of IDIR<sub>(1)</sub> at *c\** will always be an accurate, literal report what Mary said by uttering (1) at *c*. This indicates that there is a context-invariant content expressed by all utterances of (1) that can be specified disquotationally: every utterance of ‘Bill is ready’ expresses the minimal proposition *that Bill is ready*. This belies the contextualist claim that certain features peculiar to particular contexts of utterance play a constitutive role in determining the propositions expressed by utterances of (1).

As for IDIR<sub>(2)</sub>, however, the occurrence of the standard indexical ‘I’ brings it about that such a report might be false for some choice of *c* and *c\**, as we cannot be sure that the expression will pick out the same individual at both contexts of utterance. This goes to show that sentences that do not express the same minimal proposition across all contexts are those that contain standard indexicals (C&L 2005, p.143).

The *collective description* (CD) test shows that ‘context sensitive expressions block collective descriptions’ (C&L 2005, p.99). To explain, imagine that (2) and (4) are uttered at two distinct contexts:

(3) Herman is here

(4) Ernie is here

One cannot infer from the truth of these utterances that an utterance of the following collective description will also be true:

CD<sub>(3&4)</sub> Herman and Ernie are here

The expression ‘here’ is thus identified as context sensitive. Note, however, that if the context sensitive predicate were to be replaced by a context *insensitive* predicate, e.g. ‘is male’, one *could* then validly infer *that Herman and Ernie are male*.

Finally, the *inter-contextual disquotational* (ID) test shows that an utterance *S* is genuinely context sensitive only if there ‘are (or can be) false utterances of ‘*S*’ even though *S*’ (C&L 2005, p.105). So:

(5) These are fancy

(6) Cats chase rats

If we submit (5) to the ICD test, we get the (potentially) true sentence:

ICD<sub>(5)</sub> ‘These are fancy’ is false even though these are fancy

It is thereby shown that ‘these’ is a context sensitive expression, as it need not pick out same set of individuals at each context. If, alternatively, we submit (6) to the test, we intuitively end up with a contradiction:

ICD<sub>(6)</sub> ‘Cats chase rats’ is false even though cats chase rats

So, neither ‘cats’, ‘chase’, nor ‘rats’ belong to the basic set, and (6) expresses the same minimal proposition at all contexts of utterance.



As we can see, then, C&L's tests do indeed serve to demarcate a basic set of context sensitive expressions, which turns out to be the same set of expressions specified in Kaplan's list of standard indexicals (Kaplan, 1989); but whereas Kaplan never gave any definite arguments for why the phenomenon of context sensitivity should be thought to be limited to this select set of expressions, C&L deem the tests to enforce a robust limit on the extent to which context sensitivity needs to be accounted for in semantics. In fact, C&L regard the tests to be so in line with our intuitive grasp of context sensitivity that '[n]o semantic theory should classify as context sensitive any expression that does not pass these tests' (C&L 2005, p.151).

Most importantly, the tests show that sentences that do not partly comprise standard indexicals express the same invariant content over all contexts of utterance. The semantic content of any such sentence is a minimal proposition, which is the 'content that can be grasped and reported by someone who is ignorant about the relevant characteristics of the context in which an utterance of *S* took place' (C&L 2005, p.143). Hence, given that minimal propositions can be specified purely by means of disquotation, a Tarski-style truth theory, which issues in disquotational specifications of truth-conditions for sentences that do not contain standard indexicals, will, as Davidson proposed, serve as a theory of meaning.

Now, of course, minimal propositions are not the only propositions, or even the most interesting or relevant propositions, such sentences are used to express. So, C&L argue that we should be pluralists about speech act content, and acknowledge that there are a number of propositions we can succeed in communicating with our utterances, many of which can only be determined by exploiting our knowledge of the contexts at which such utterances take place. Of all the propositions that might be communicated by an utterance, though, the minimal proposition is that unique proposition that is compositionally determined by the meanings of a sentence's constituent expressions and their syntactic arrangement. On C&L's view, that there is such a proposition *semantically* expressed, which is (almost) completely immune to the various effects context can have on the propositions utterances *intuitively* express, must be presupposed

in order to explain how speakers are able to communicate successfully across diverse contexts:

The proposition semantically expressed is that content the speaker can expect the audience to grasp (and expect the audience to expect the speaker to expect them to grasp) even if they have mistaken or incomplete communication-relevant information....the proposition semantically expressed is our minimal defense [*sic.*] against confusion, misunderstanding, mistakes and it is that which guarantees communication across contexts of utterance. (C&L 205, p.184-185)

Ultimately, then, C&L take semantic minimalism paired with a theory of *speech act pluralism* not only to provide good, intuitive evidence for the claim that sentences express minimal propositions, but also to provide a plausible explanation of the cognitive role of such propositions; i.e. minimal propositions play the role of a kind of ‘shared fallback content’ (C&L 2005, p.185) that is the content we can reliably grasp exclusively on the basis of our semantic competence.

### 1.3. Borg’s semantic minimalism

Borg (2004, 2012) provides a thorough defence of the semantic minimalist claim that ‘there is a level of propositional or truth-evaluable content, namely sentence meaning, which can be delivered through interpretation of the formal features of the expressions in play’ (Borg 2004, p.74). In adopting this view, Borg rejects contextualist claims that utterance interpretation involves taking into account various pieces of information about contexts of utterance (i.e. information about, say, speakers’ intentions) that are not encoded in the lexico-syntactic features of linguistic expressions.

An important motivation for Borg’s defence of semantic minimalism is the idea that there are good reasons for thinking that our semantic competence, i.e. the competence that underpins our ability to interpret the semantic contents of utterances, is modular (in the sense of Fodor (1983)). That is, the kinds of processes involved in recovering the semantic contents expressed by linguistic utterances exhibit all the hallmarks of modular processes; e.g. they are fast, unconscious, automatic process that are informationally encapsulated in that they operate over a highly select set of domain specific inputs (for

an extensive defence of this, see Borg 2007, pp.86-106). Borg emphasises that if we accept the claim that semantic competence is modular, we must thereby accept that the recovery of semantic content proceeds via the implementation of a range of purely deductive, computational operations over the formal, lexico-syntactic features of sentences. A semantic theory aimed at giving an adequate account of semantic competence will thus take the form of something like a Tarski-style truth theory that specifies the semantic contents of sentences exclusively in terms of their lexico-syntactic features.

Importantly, however, what we cannot expect from any such theory is that it pair sentences with the contents speakers intuitively take them to express at particular occasions of use, since figuring out the intuitive speech act contents of utterances typically involves implementing rich abductive, extra-linguistic processes (yielding inferences from various kinds of contextual information, such as speakers intentions, to what is said) (Borg 2012, p.12). So, assuming ‘a Fodorian, modular story about our semantic competence’, we should endorse the minimalist claim that our linguistic understanding can be accounted for ‘in terms of discrete, deductive, syntax driven reasoning processes, rather than via the murky world of content-driven inference to the best explanation’ (Borg 2012, p.13).

According to Borg, then, the processes involved in determining semantic content cannot ‘look beyond’ the formal features of sentences, as this would conflict with the claim that such processes are modular, and so informationally encapsulated.

Borg further argues that we can maintain this view even when it comes to explaining how speakers determine the semantic contents of sentences containing standard indexical and demonstrative expressions, despite the fact that identifying the referents of such expressions often appears to involve taking into account extra-linguistic perceptual information, information about speakers’ intentions, and so on. To show how, Borg begins with an appeal to Higginbotham’s (1994) idea that we can specify the truth-conditions of sentences in which, say, demonstratives occur by using *conditionalized* T-sentences of the following form:

If the speaker of 'this is red' refers with the utterance of 'this' therein to  $x$  and to nothing else, then that utterance is true iff  $x$  is red. (Higginbotham 1994, p.93)

As we have seen, the antecedent of such conditionals specifies a condition for identifying the referent of the demonstrative expression in terms of its character, whereas the consequent specifies the truth-conditions of the sentence for a referent thus identified at some context. The information pertaining to the character of the demonstrative forms part of our knowledge of the lexical entry for 'that', placing constraints on how we identify its referent, but it makes no appearance in the truth-conditions of the utterance, which are specified in the standard way.

Developing this, Borg proposes that when one encounters an utterance of, say, 'That is mine', a formal representation of the semantic content of the sentence uttered is generated purely on the basis of the lexico-syntactic properties of the sentence. The representation is of the form  $\alpha$  is  $\beta$ 's, where  $\alpha$  and  $\beta$  are syntactically generated *singular concepts* that are created for the expressions 'that' and 'mine' respectively, and whose semantic contents are the objects to which they refer. At this point, someone who knows the language knows what conditions would have to obtain for the utterance to be true, and it follows that an utterance of 'That is mine' is thus true iff the object referred to by the singular concept  $\alpha$  created for 'that' belongs to the object referred to by the singular concept  $\beta$  created for 'mine'.

That said, one still has to identify the objects to which  $\alpha$  and  $\beta$  refer in order to determine *whether* the sentence is true. Doing this involves utilising one's antecedent knowledge of the characters of the relevant expressions, so as to fix what *types* of objects to which the singular concepts  $\alpha$  and  $\beta$  refer, as well as, e.g., one's perceptual knowledge of the surroundings at which the utterance takes place. Again, the conditions for identifying the referents of  $\alpha$  and  $\beta$  are not specified as part of the sentence's truth-conditions:

If  $t$  is a token of 'that is mine' uttered by  $\beta$ , and the token of 'that' therein refers to  $\alpha$  then  $t$  is true iff  $\alpha$  is  $\beta$ 's. (Borg 2004, p.206)

On Borg's view, then, we can specify the truth-conditions sentences containing standard indexicals and demonstratives exclusively in terms of their lexico-syntactic features despite the fact that identifying the referents of such expressions involves implementing process that are strictly speaking extra-semantic. She goes on to emphasise the following:

The important point here is that there is a difference between grasping the truth condition for an utterance and being able to verify whether or not that condition is met. In the case of novel uses of demonstratives and indexicals, given linguistic information alone, an agent is capable of grasping the truth conditions of utterances containing the terms, though she is not yet in a position to see that they are satisfied (she knows *what it would be* for the utterance to be true but she doesn't yet know that it *is* true). Now it seems to me that this should be sufficient for counting as grasping the literal meaning of the sentence produced (that is, as grasping genuinely semantic content), so I reject the idea that our minimal semantics for demonstratives and indexicals is too minimal to count as truly semantic. (Borg 2004, p.165)

Another interesting aspect of Borg's approach that will be relevant in what follows is that she deviates from C&L's view that minimal propositions can always be specified purely disquotationally. For Borg, maintaining a modular, hence minimalist, conception of semantic content does not necessarily involve ruling out speakers' intuitions about incompleteness. So, to return to our earlier example, semantic minimalism can accommodate speakers' intuitions to the effect that in order for utterances of the sentence 'John is ready' to be true, it must be the case that John is ready *for something*.<sup>38</sup> The important point is that this claim does not depend on intuitions about 'one-off' effects particular contexts have on the content expressed by utterances of the sentence, but 'type-level' intuitions about the context-independent meaning of 'ready'.

---

<sup>38</sup> This conflicts with C&L's view that utterances of 'John is ready' do *not* express the proposition *that John is ready for something or other*, but rather that they all express the proposition *that John is ready...*' (C&L 2005, p.97).

Borg (2012, pp.100-102) argues that we can accept that there is something about the semantics of 'ready' that gives rise to the intuition that all utterances of 'Bill is ready' express the minimal proposition *that Bill is ready for something*. This does not belie the modular approach as it involves an appeal to lexically encoded information that forms part of our semantic competence. So, whilst our semantic competence alone does not enable us to determine the speech act content we associate with utterances of the pertinent sentence, e.g. *that Bill is ready for an evening out*, it does enable us to determine what Borg (2004) refers to as the *liberal truth-conditions* of the sentence, which are somewhat richer than its purely disquotational truth-conditions. Borg goes on to present an analysis of why it is that certain expressions bring about incompleteness effects, but I shall reserve further discussion of her views on that matter for the moment. For now, I take it enough has been said to show why Borg thinks we should defend semantic minimalism, why semantic minimalism takes itself to succeed in defending the prospects of developing truth-conditional semantic theories, and what kinds of factors enter into the determination of minimal propositions.

## 2. Objections to semantic minimalism

Before looking in detail at the potential problems facing the two individual accounts of semantic minimalism analysed above, some observations on the overall argument form and strategy deployed by the semantic minimalists generally are in order.

First, the semantic minimalists polarise the situation in the following way. They readily acknowledge the potentially serious destabilising consequences of context, and regard the attempts of the mild contextualists to minimise its effect as a vain attempt to halt the inevitable. They also claim that such radical disruption, were it to apply to semantic content, would necessarily lead to the loss of the systematicity that they consider essential in preventing our success in communication from being simply mysterious. These factors licence, in their view, the separation of speech act content from semantic content, in the way that we have seen.

The argument thus appears to hinge on the idea that the effects radical contextualism inevitably lead to a lack of *systematicity*, whereas an equally valid interpretation of their

words is that what is at stake is, in fact, the issue of *tractability*. That is, it *may* be the case that the method speakers use to interpret sentences relative to contexts of utterance is systematic even though it operates on information pertaining to the formal features of sentences *and* information pertaining to the contexts in which those sentences are used. Assuming that this, admittedly speculative, proposal is at least plausible, it follows that that what minimalist arguments actually show is that a method of interpretation that operates with contextual information will not be tractable from the perspective of a standard truth-conditional semantic theory.

Put somewhat differently, if the problem minimalists raise ultimately comes down to a problem of tractability, it leaves open the thought that, although radical, contextual effects might nevertheless be systematic, either in terms of the operation of the wider cognitive processes or, to some degree at least, those within the language faculty (the latter approach is discussed in the next chapter). To be sure, such systematicity may be extremely complicated (although see shortly) and, perhaps, currently intractable to human enquiry, but its very possibility undermines the minimalists' warrant for the necessity of the radical reduction in the scope of semantic content that they propose unless the theoretical objects of linguistic analysis are to be determined by what is convenient to theorists. Of course, tractability difficulties might argue for a restricted focus on methodological grounds, but not in and of themselves for the positing of radically different notions of content.

Second, and following from the above, as part of the way in which they set up the argument, the semantic minimalists seem to offer an exhaustive disjunction in describing the phenomenon of communication: it is either ideal, systematic, and explicable, or it is ideal, unsystematic and mysterious, and we are invited, as we have seen somewhat rhetorically, to choose the former, there being no other options. However, it is entirely plausible that communication might just be 'good enough', rather than ideal. Although the process appears very impressive in operation, and the exhibition of such proficiency invites us to infer some kind of idealised underlying mechanism (this is a familiar notion in prominent evolutionary psychology arguments (e.g. Cosmides and Tooby 1995), the actual system that delivers the function may be

relatively simple and so well adapted to the environment in which it is deployed that its success rate is extremely high. For instance, the mechanism subserving communication might, in part at least, be driven by heuristics in the way that Gigerenzer (2000) and others propose for other key components of our cognitive repertoire. Adopting such an approach would offer the potentially interesting consequence of denying that communication needs to be as precise as we might presume, and might also gesture at an account of how we seem initially to quickly assume particular contexts, and therefore interpretations, and are then occasionally proved wrong. I will further contrast the properties of such systems with Fodor's (and Borg's) characterisation of abductive inference in analysing problems with Borg's specific account shortly.

However, for present purposes, the key consequence of this possibility is to weaken the pressure on us to accept the element of the disjunction that the minimalists wish us to choose, the above offering the further option (with some evolutionary precedent and cognitive plausibility) that communication may be tractable to enquiry without our having to make the move the minimalists urge. Thus, whilst the minimalists' argument is strictly an empirical and rhetorical one (we are invited to choose one element of the disjunction because the other is painted as preposterous), their strategy of neglecting other plausible alternatives gives the impression that their argumentation has greater modal force than it actually warrants. Seeing the 'first move in the conjuring trick' as questionable is key because, as they set it up, once one has chosen the minimalists' 'systematic and explicable' branch, some brand of minimalism seems to follow almost deductively. However, recognising that there are, in fact, a number of other alternatives available radically reduces the force of this core argument, and again raises questions regarding the support for their radical postulations regarding semantic content.

I will return to generic objections to the minimalist position at the end of the chapter, but in the next two subsections I will examine the key arguments of their two main proponents, C&L, and Borg.



## 2.1. Objections to C&L

I shall raise two main objections that apply specifically to C&L's version of semantic minimalism. First, I shall raise some problems for their claim that disquotational tests serve to delimit the set of context sensitive expressions and to identify minimal semantic content. Second, I look in more detail at some issues pertaining to the claim that the appeal to disquotational reportage provides good evidence for the claim that sentences express minimal propositions.

Do C&L's disquotational tests show that contextualists have made an empirical error about the extent to which natural languages are context sensitive? I think there are good reasons for arguing that, in and of themselves, they do not. One initial problem for the idea that disquotational tests serve adequately to delimit the range of context sensitive expressions is that numerous expressions ranging beyond C&L's basic set actually pass the tests that C&L say they fail. For example, to use an argument from Szabó (2006, pp.35-6), imagine that Mary utters (7) in the context of a discussion about Sesame Street characters:

(7) Ernie lives on Sesame Street

Let us then imagine that the following IDIR is produced for (7) in the context of a discussion about the philosopher, Ernie Lepore:

IDIR<sub>(7)</sub> Mary said Ernie lives on Sesame Street

This report of what Mary said is obviously false, so the proper name 'Ernie' passes the test for context sensitive expressions, as well as the other tests C&L set forth.

Similar observations have been raised for predicates such as 'ready' (Recanati 2006), 'weighs 80kg' (Bezuidenhout, 2006), 'grunts' (Travis, 2006), and so on. To take just the first of these examples, one may well argue that an IDIR of the form 'Mary said that Bill is ready' can be false, as it might fail to report what Mary said when she uttered the relevant sentence in the relevant target context. In other words, although Mary might have uttered the original sentence in the context of a discussion about whether Bill is ready to be promoted, the report might be made in the context of a discussion about

whether it is time for Bill to retire, and so this would then be considered a false report of what Mary actually said.

The above points seem to show that disquotational tests, in and of themselves, do not effectively resolve the kinds of issues raised by CSAs, and a committed contextualist will argue that CSAs can be employed in such a way as to influence speakers' intuitive judgments about what it is such tests actually show.

As a result, if one is committed to the view that determining the truth-conditional content of an utterance of 'Bill is ready' will involve taking into account what Bill activity Bill is ready to perform or engage in at a given context, then this will influence one's judgments about whether a disquotational report of an utterance of the sentence is a true report of what a given speaker said by performing such an utterance. Consequently, disquotational tests only serve to delimit the set of context sensitive expressions if one is already sceptical about the import of CSAs.

One might nevertheless argue that there is something particularly striking about the kinds of effects that occurrences of standard indexical expressions have on disquotational reportage. Indeed, it may well be that such expressions fail disquotational tests in a particularly transparent, obvious manner. However, this would seem merely to show that there is a qualitative difference between the semantic behaviour intuitively exhibited by standard indexicals and the semantic behaviour of other expressions. Such a result would be insufficient to establish the claim that standard indexicality is the only legitimate variety of context sensitivity, and so would not support the conclusion that the distinction C&L's tests establish is *the only* distinction to be made between context sensitive and context insensitive expressions. As Recanati writes:

I do not want to deny that there are differences between expressions in the Basic Set and other (alleged) context-sensitive expressions, such as 'ready'; nor that some of those differences may have to do with the availability of homophonic reports, collective descriptions, etc. But the conclusion should not be that expressions like 'ready' are not really context sensitive. Rather, one should be

prepared to acknowledge differences within the class of context-sensitive expressions: they don't all behave similarly. (Recanati, 2006, p.29)

Even within the basic set, Recanati continues, there are differences between the way that C&L's so-called genuinely context sensitive expressions behave. For example, 'now' (as opposed to, say, 'here') *can* refer to the same time it was originally used to refer to even when disquotationally reported at a later time (and so does not always block the IDIR test). That is:

I can report John's utterance of 'Now I am hungry' by uttering, much later, 'He said that now he was hungry'. One should look at all this in much more detail than C&L themselves are willing to do. The more we extend the list of context-sensitive expressions, the more we should be disposed to register, and explore, fine differences within the class. (*ibid*)

The points raised here have serious consequences for C&L's claim that disquotational tests serve to delimit the class of context sensitive expressions. What is more, in showing that disquotational tests in and of themselves do not provide an absolute distinction between the context sensitive and the context insensitive expressions of a language, the points made here raise questions for the view that any sentence that passes these tests expresses a context invariant minimal proposition.

In response to such attacks on their tests for context sensitivity, C&L adopt the following line:

[In *IS*] we never claim passing these tests is either necessary or sufficient for context sensitivity: We describe the tests as providing 'evidence' that sentences or expressions are context sensitive. (C&L, 2006, p.51)

The first problem here is that the above statement conflicts with C&L's earlier claims. Recall that C&L explicitly state that '[n]o semantic theory should classify as context sensitive any expression that does not pass these tests' (C&L 2005, p.151). Second, even if we set aside this inconsistency and take C&L to be here opting for a more modest view of their tests' effectiveness, there remains the problem that without a definitive

method for identifying context sensitivity (which also demonstrates that it is severely limited), they appear to have little warrant either for assuming that one can rule out the presence of genuinely context sensitive expressions in most sentences, or for claiming that most sentences thus express minimal propositions. I take it then that the objections raised here thus present a real problem for C&L's brand of semantic minimalism.

Moving on to the second point, C&L argue that the fact that we are able to produce true disquotational reports of utterances of sentences across diverse contexts ensures that such sentences semantically express a minimal proposition. That is, our ability to produce an accurate disquotational report of an utterance is taken to depend on our ability to grasp the proposition semantically expressed by that utterance. In short, reported content is semantic content (C&L 2005, p.152).

Adopting this line of argument, C&L claim that even if we have no intuitive grasp of what would need to be the case in order for it to be true that an individual is ready – i.e. even if we intuitively take it that 'ready' does not denote a property – the fact that we can produce a disquotational report of a sentence such as (8) is enough to ensure that the sentence has a minimal semantic content:

(8) Bill is ready

The problem with this kind of argument, though, is that it can effectively be reproduced for (9):

(9) The borogoves did gyre and gimble

Of course, speakers of English have no intuitive grasp of the kinds of conditions that would have to obtain in order for it to be true that the borogoves did gyre and gimble. That said, though, there is nothing to stop one from producing an indirect disquotational report of (9); i.e. one could say, 'Mary said the borogoves did gyre and gimble'. The problem is that no one would take it that, in this case, the mere fact that we can produce an indirect disquotational report for (9) ensures that the sentence expresses a minimal proposition.

Now, one immediate response to this objection is simply to make a distinction between use and mention. That is, one might argue that we do not actually *use* (9) when we report what a person who uttered that particular string of sounds said. Rather, we merely *mention* (9), and so what we actually produce when we report utterances of the sentence are direct rather than indirect reports of those utterances. The question, though, is why the same cannot be said for (8). That is, if we have no intuitive grasp of what it would be for an individual to be ready *simpliciter*, then why do we assume that the mere fact that we can produce a report of (8) ensures that ‘ready’ makes a determinate contribution to the truth-conditional content of the sentence? The only explanation seems to be that we are simply assuming that ‘ready’, unlike the nonsense words that make up (9), makes a determinate contribution to truth-conditional content. But this is precisely the issue that the appeal to disquotational reportage was supposed to settle.

Ultimately, the point is that the reason we reject the claim that we can produce disquotational reports of utterances of (9) is that (9) does not express anything like a complete proposition or determinate set of truth-conditions. But then this shows that we simply fall back on our intuitive grasp of which sentences do and which sentences do not express truth-conditional content when we come to decide which sentences we can and which sentences we cannot produce disquotational reports for. So, it seems that the possibility of disquotational reportage is not, in and of itself, a good indicator of whether sentences express determinate sets of truth-conditions.

## 2.2. Objections to Borg

In this section I will question Borg’s account from two perspectives. The first will focus on Borg’s presumption is regarding mental architecture, and build on issues raised earlier concerning the way in which minimalists couch the problem of the apparent systematicity in communication. The second will look at the possible nature of minimal content, making some comparisons with the narrow content of Fodor.

In the earlier analysis of Borg’s position, we saw that two inputs were key to her strategy. The first was a somewhat rhetorical argument based on an exhaustive disjunction, which invited us to accept semantic minimalism and its attendant changes in

semantic content on pain of mystery. The second was the adoption of a Fodorian conception of the architecture of cognition, along with his idea of the way which abductive reasoning works, and the associated arguments that he deploys. From these ingredients we were intended to conclude, with Borg, that only an extremely restricted sense of content was plausible, given the modular nature of the language faculty and the constraints on its operation in virtue of its domain specificity and informational encapsulation.

I discussed the problems associated with the first of these inputs at the start of this section, highlighting the confusion between systematicity and tractability, along with at least one addition to the list of options proposed by the exhaustive disjunction. These considerations, I argued, materially weakened Borg's case. I now want to look in more detail at her architectural assumptions for cognition and the language faculty, as well as her reliance on Fodor's argumentation.

Borg's invitation us to assume that the mechanism subserving the process of recovery of semantic information exhibits the hallmarks of modular systems is intended to show that the scope of semantic content must be restricted. There are two main reasons offered for presuming modularity, both of them are based on Fodor's architecture, and both are somewhat consonant with the semantic minimalists' more general picture portrayed by the exhaustive disjunction discussed earlier. In other words, both minimalists and Fodor view pragmatic, abductive processes as both mysterious and potentially unbounded, whereas modular processes are strictly controlled and clearly systematic.

The attractions of the Fodorian system for Borg are therefore obvious. However, the presumption of a modular architecture for cognition more generally is contentious, even if it is more generally accepted that modularity for the language faculty is plausible. Many evolutionary psychologists, computational biologists and philosophers of mind (Pinker, Gallistel, Cosmides and Toobey, Carruthers, etc.) broadly advance a competing theory: namely, the idea that the mind is almost completely (or massively) massively modular. If this architecture were presumed rather than Fodor's extremely restrictive deployment of modules, the argument from modularity to semantic minimalism would

not go through – evidence of modularity would not restrict the possible domain of operation for semantic recovery to a single module (the language faculty).

In fact, Borg does consider the case for massive modularity (2004, pp. 78-80), but rejects it, yet again following Fodor's line of argumentation, or almost, concluding that massive modularity cannot cope with global and abductive processes. It is interesting, however, that Fodor considers the case of massive modularity specifically because he entertains the idea that it might actually offer a solution to the inherent and well-known problem associated with thought more generally and global abductive reasoning in particular: namely, the frame problem. He considers how heuristics, as modular processes, might work, but concludes against their deployment on the grounds that he cannot see how the appropriate heuristic could be selected without some sort of homuncular supervisor. He dubs his *a priori* argument to this end 'the input problem'.

The problem for Borg is that this argument bears all the weight in Fodor's (2000) exposition, and it is easily defeated. Consideration of the 'production' architectures, for example, reveals a possible way in which selection could be achieved bottom-up rather than top-down, and, this being, importantly, an *a priori* argument, the mere possibility is sufficient to undermine Fodor's (and therefore Borg's) position.

A further criticism levelled at Fodor in this context is that he somewhat idealises the capability of human abductive reasoning in practical circumstances. Pinker (2005, p.10) points out that when Fodor talks about human abduction, he spends most of his time discussing examples from the history of science. He, Fodor, alludes, for example, to the arguments of Quine and Duhem concerning entire systems of belief, as if such examples were representative of the everyday abductive processes human beings are actually involved in. Once the disparity between scientific reasoning on the one hand, and everyday reasoning on the other, is recognised, the very way in which he has framed the abduction problem is called into question. As a result, the credentials of the alternative, highly restricted processes for accomplishing such evidence, such as "fast and frugal" heuristics (Gigerenzer et al. 1999, p.20) appear more plausible.

These specific findings concerning Fodor's (and Borg's) conceptions of human reasoning processes, and their associated architecture, serve further to undermine the reasons posited by Borg for the proposed radical change in semantic content. It turns out that pragmatic processes may well indeed be systematic, albeit heuristic, exhibiting the very advantages (fast, automatic, systematic) that Borg presumed based on her adoption of Fodor's theories of cognitive architecture were exclusive to the language faculty. Whilst this may not be fatal to her cause, given the strength of her reliance on this cognitive picture it surely represents a severe weakening of her stance.

I now examine the nature of minimal content postulated by Borg in an attempt to articulate what such content might consist in, and whether its postulation is consistent with an explicitly externalist semantics. The locus of this analysis is Borg's claim examined earlier that there is something about the semantics of 'ready' that gives rise to the intuition that all utterances of 'Bill is ready' express the minimal proposition *that Bill is ready for something*. This is a weaker claim than C&L, but nevertheless I think there are criticisms that can be levelled against it.

The principal concern is that the content involved is vacuous. If this charge were upheld, it would deprive minimal content of any participation in externalist semantic determination. The chief problem is that the property of 'being ready *for something*' seems to be one that all objects always possess. Dogs, for instance, are always ready to be one of, and perhaps all of, the following: walked, put to bed, cooled or heated, fed, destroyed in a nuclear blast, eaten by an indiscriminate wild animal etc. Even atoms are constantly ready to be bound, observed and split. If we therefore say that 'Bill is ready', and in doing so express the proposition *that Bill is ready for something*, we appear to have said something that is always true in some sense. We also appear to have said something that does not distinguish Bill's state from that of any other object. Clearly, the application of a particular context would create conditions under which the proposition could be false, and Bill's state could be distinguished from that of other objects, but merely knowing that Bill is *ready for something* does neither, and this would seem to be fatal for the employment of minimal content in truth conditional semantics.



The above would seem to argue that pragmatics, or at least some process of wider cognition, is necessary for the determination of content. Given our earlier discussion of the possible nature of wider cognitive processes (contra Fodor and Borg), this need not give cause for concern, unless one happens to be a semantic minimalist. Nevertheless, it does also open up the possibility that what the analysis of Borg's position exposes is that the conceptual system outside the language faculty may possess resources that are finer grained than that of the lexicon. In other words, whilst it is presumed that 'readiness' as a lexical item is taken to have a mirroring counterpart in the conceptual system, this might well be false. Instead, the individuation of the conceptual system may be far more fine-grained, thereby distinguishing, for instance, between cases where the subject is psychologically ready to initiate an action, and those where he or she (or it) is simply in an appropriate physical state to enable the actions of others. This is clearly speculation, but such a system would enable a more fine-grained, unsystematic, distinction between cases of readiness and demonstrate how the building blocks of an externalist theory of content might be formed. Such a possibility would, of course, be of no comfort to Borg.

One last comment is perhaps in order on the notion of minimalist content, and this reintroduces Fodor's theories. Given Borg's reliance on Fodor in other areas, it is striking how close his account of narrow content is to hers of minimal content. I do not intend to contribute a 'blow-by-blow' comparison, and Fodor is talking about mental states where Borg is talking about sentences, but it is illustrative of the similarities between the two positions that for Fodor, narrow content represents 'a notion of some mental state that generalises over the environments relative to which mental states are semantically evaluable' (Fodor 1991, p.300). In other words, narrow content is the consistent core that remains when a particular mental state is abstracted from any of its possible contexts. And in a similar way to Borg's 'Bill is ready for something', narrow content is waiting to be anchored in a particular context where its truth evaluable content will be settled.

Fodor is clear however, that narrow content itself is not truth evaluable, and that if one requires content to have this property, then the intrinsic content that these mental states share is not, strictly, content at all, since the mental states only possess this type of

content ‘potentially’. Whilst the parallel with Borg is clearly not exact, it is interesting that Fodor considers context to be necessary for broad content.

### 3. Summary

It seems, then, that there are a number of powerful arguments against the minimalist position, based on their argument form, assumptions and conclusions.

First, we have seen that the way in which the minimalists set up the argument are largely rhetorical and based on a particular, and disputable, conception of the way in which cognition operates outside the language faculty. In the case of Borg, reliance on Fodor’s position to establish modularity for the language faculty and some form of mysterious abductive reasoning processes for the rest of cognition will only be as strong as Fodor’s arguments for his position. These turn out to be relatively weak, and we are left with the possibility that Borg’s inability to conceive of systematicity occurring in cognition outside the language faculty has forced to adopt an unnecessarily restrictive notion of semantic content, that may, in fact, be no content at all.

Second, it turns out that the cross contextual claim, that funds much of C&L’s position is not all that it seems. In particular, the crucial direct/indirect distinction simply fails to establish their claims, as we have seen – it turns out that there are no grounds that are not question begging for determining whether a direct or indirect communication is being received. This is a substantial problem as it undermines a significant element of the cross contextual claim.

Now an important point to conclude on is that I am not at all opposed to the basic notion, which provides much of the motivation for semantic minimalism, that when we come to investigate how the formal properties of *sentences* contribute to interpretation, we should be very wary of the idea that extra-linguistic effects will have anything more than a very minimal bearing on that investigation. That is, I too think we should adopt a minimalist perspective on the extent to which extra-linguistic factors are relevant to the interpretation of language. That said, though, what I hope to have shown in this chapter is that the cost of *going minimalist* may well be that we no longer get to assume that interpretations of *sentences* will be interpretations of truth-conditions.

## CHAPTER 4

### THE APHONIC APPROACH

The purpose of the chapter is to provide a critical assessment of the *aphonic approach* set forth by Stanley (2007).<sup>39</sup> The approach is predicated on the claim that ‘all truth conditional effects of extra-linguistic context can be traced back to logical form’ (Stanley, 2007, p.30). Stanley’s chief aim is to show that the logical forms of various sentences partly comprise *aphonic pronominal elements*, which are phonologically null constituents of sentential structure that have their semantic values assigned to them relative to contexts. In adopting this view, Stanley rejects the minimalist claim that all genuine contextual effects on content can be traced back to occurrences of standard indexical expressions. What is more, in arguing that there is a variety of elements situated in the logical forms of sentences whose semantic values are assigned relative to certain specific features of the contexts in which those sentence are used, Stanley is led to argue that ‘there is no gap between the linguistically determined content of a sentence, relative to a context, and the proposition it intuitively seems to express’ (Stanley 2007, p.5).

The chapter proceeds as follows, and is in two parts. First, I offer an exposition of the aphonic approach, focusing on the arguments Stanley sets forth in order to defend the claim that there are aphonic pronominal elements situated in the logical forms of sentences. Second, I raise a number of challenges for these arguments, and offer an alternative analysis of some of the more compelling examples Stanley relies on in order to justify the postulation of aphonic structure. I conclude that it is difficult to see why the appeal to aphonic elements would help explain why context sensitive expressions appear to be systematically associated with certain gaps in truth-conditional content.

---

<sup>39</sup> The term *aphonic approach* comes from Pupa and Troseth (2011).

## 1. The argument from binding<sup>40</sup>

In order to defend the claim that there are aponic elements situated in the logical forms of sentences, Stanley sets forth the *argument from binding*. The argument proceeds in three stages: first, a variety of examples are provided to establish the claim that the ‘gaps’ or ‘slots’ that speakers intuitively associate with occurrences of context sensitive expressions, which are typically saturated relative to context, can actually be bound when those expressions occur in the domain of a binding operator; second, it is shown that contextualist accounts, which Stanley refers to as *pragmatic approaches*, cannot account for the relevant instances of binding; third, it is argued that instances of semantic binding only occur when there is a corresponding instance of syntactic binding. This final point is taken to provide support for the claim that the gaps speakers intuitively associate with context sensitive expressions actually correspond to aponic elements that are located in the syntactic structures of those sentences.

To assess the aponic approach, I begin by examining the *binding examples* Stanley sets out. I then look at the arguments Stanley presents in order to show that pragmatic approaches are unable to account for the relevant examples. Finally, I give an outline of Stanley’s attempt to account for the kinds of interpretive phenomena binding examples serve to reveal.

### 1.1. Context sensitivity and binding

We have seen that when a context sensitive expression occurs in a sentence, its presence gives rise to an intuitive *gap* in truth-conditional or propositional content. Speakers saturate such gaps by taking into account certain specific pieces of information gleaned from contexts of utterance. So, in the case of (1a), we have a gap that is filled by

---

<sup>40</sup> In order to keep things focused, I will limit my attention to Stanley’s argument from binding, which I take to be his primary argument for the presence of aponic elements at LF. Stanley also sets forth an argument from weak crossover effects, and an argument from strict-sloppy readings. Pupa and Troseth (2011) set forth a highly detailed, and highly critical analysis of these arguments.

information pertaining to the contextually salient domain to which the quantifier ‘three students’ is restricted (as in (1b)); and in the case of (1c), we get a gap that is filled by information pertaining to a contextually salient location (as in (1d)):

- (1) a. John failed three students
- b. John failed three students *<in John’s class>*
- c. It is raining
- d. It is raining *<in Paris>*

Stanley describes a range of binding examples that serve to show that the kinds of gaps that arise in instances of context sensitivity need not always be filled from context. That is, when such gaps occur within the domain of an appropriate binding operator (e.g. a quantifier phrase), speakers are able to get bound readings of the relevant gaps. In order to construct such examples for (1a) and (1c), all we need to do is embed each of those sentences within a more complex linguistic setting where the gaps associated with the context sensitive expressions ‘three students’ and ‘rain’ fall within the domain of, and are bound by, a quantifier phrase that occupies a position *higher up* in sentential structure. Modifying the original sentence slightly in order to do this, the resulting sentences we are left with are (2a) and (2c), whose interpretations are paraphrased in (2b) and (2d):

- (2) a. In most of John’s classes, he fails three students
- b. In [most  $x$ :  $x$  is a class of John’s] he fails three students *<in  $x$ >*
- c. Everywhere I go, it rains
- d. [Every  $x$ :  $x$  is where I go], it rains *<in  $x$ >*

Similar effects occur in relation to the following sentences:

- (3) a. Most species have members that are small
- b. [Most  $x$ :  $x$  is a species]  $x$  has members that are small *<for  $x$ >*
- c. Every student answered every question
- d. [Every  $x$ :  $x$  is a student]  $x$  answered every question *< $x$  was asked>*

Stanley takes it that these examples serve to reveal certain semantically relevant facts about the way speakers interpret utterance of sentences that contain context sensitive expressions. Hence, they are facts that a semantic theory ought to attempt to account for. Before going on to look at how different approaches fare when it comes to the task of accounting for the relevant facts, I want to point out an important difference in the way these examples operate.

I shall refer to the examples listed in (2) as *strong binding examples*. Such examples reveal particularly striking constraints on the ways speakers are able to interpret utterances of the relevant sentences. That is, it seems that the interpretations of (2a) and (2c) that were offered in (2b) and (2d) are the *only* interpretations of truth-conditional content that speakers are able to intuitively associate with utterances of those sentences. In short, then, speakers *must* interpret (2a) and (2c) as involving some form of binding over an implicit gap in truth-conditional content. The reason for this is that it seems that if we do not interpret these sentences as involving binding over the implicit gaps, then we effectively end up with instances of *vacuous binding*, which occurs when a binding operator, such as a quantifier phrase, does not actually bind any variable. Stanley quite right argues that instances of vacuous binding signal a failure of interpretation (see Stanley 2007, p.195).

I refer to the examples of (3) as *weak binding examples*. In each of these cases, it appears that speakers can interpret utterances of the relevant sentences in a way that does not involve any instance of binding over the implicit gaps associated with the context sensitive expressions that occur in the sentences. That is, one might take it that (3a) and (3c) could be assigned the interpretations paraphrased in (4) relative to some context:

- (4) a. [Most  $x$ :  $x$  is a species]  $x$  has members that are small *<compared to an elephant>*<sup>41</sup>

---

<sup>41</sup> To help clarify this reading, imagine that Bill is talking to his child who has not yet seen any of the other animals at the zoo. The child is amazed by the size of the elephants, and asks “Are all animals that big?”. Bill replies: “No. Most species have

- b. [Every  $x$ :  $x$  is a student]  $x$  answered every question  $\langle$ *the contestants were asked on tonight's game show* $\rangle$

The difference between these cases and the cases of (2) is that the quantifier phrases that occur higher up in sentential structure already bind positions in the sentences they range over. That is, at LF, the quantifier ‘Most species’ quantifies over the subject position of the embedded open sentence ‘ $x$  has members that are small’. That said, there is no danger of these sentences giving rise to any form of vacuous binding, which means that speakers are free to fill those implicit gaps with information from context.

Ultimately, it will become clear in what follows that when we come to assess what kinds of semantic facts Stanley’s examples serve to reveal, it is important to keep in mind the distinction between strong binding examples, where binding of an implicit gap is mandatory for the purposes of successful interpretation, and weak binding examples, where binding over an implicit gap is optional.

## 1.2. Pragmatic approaches

Pragmatic approaches are approaches that account for instances of context sensitivity by appealing to the pragmatic competence of speakers. That is, advocates of such approaches argue that the reasons we sometimes incorporate certain pieces of contextual information into our interpretations of the truth-conditional contents of sentences are ultimately pragmatic ones. Recanati, for example, proposes that the reason speakers take it that they need to identify the location of the rain event a speaker refers to by uttering (1c) (‘It is raining’) has to do with the pragmatic fact that we ‘we care about meteorological events to the extent that we care about the locations where they take place’ (Recanati 2010, p.86). So, when a speaker utters (1c), the only way to make sense of what they are saying is to assume that their purpose is to tell us about something that is taking place *at a particular location*. As a consequence of this, we tend to take it that the location of the raining event is part of the information a speaker aims to convey

---

members that are small”, which is true just in case most species have members that are smaller than the average elephant.

when uttering the sentence, and so we take it that such utterances are true iff it is raining at a contextually salient location.

Stanley refers to such explanations of the kinds of processes speakers utilise to determine the truth-conditions of utterances as *unarticulated constituent analyses* (UCAs). An unarticulated constituent is an element of truth-conditional or propositional content that has no linguistic correlate at any level of linguistic representation. On Stanley's characterisation of these kinds of analyses, unarticulated constituents are not incorporated into interpretation to satisfy any linguistic requirements (such as the requirement that each sentential constituent must be assigned a semantic value), but instead enter into interpretation via the implementation of pragmatic processes that are not answerable to linguistic constraints.

On Stanley's view, once we allow that genuine elements of truth-conditional content can be incorporated into interpretation via the implementation of such unconstrained processes, we effectively have to give up on the idea that a semantic theory will be able to provide explanatory predictions of speakers' intuitive judgments about the truth-conditional contents that utterances of sentences express. Indeed, taken to an extreme, UCAs appear to undermine the notion that there are *any* necessary linguistic constraints on utterance interpretation (Stanley 2002a, p.161-7). That said, though, Stanley argues that UCAs can in fact be shown to be inadequate when it comes to accounting for some of his binding examples, so this gives us reason to reject accounts that appeal to UCAs.

In order to get a sense of exactly where UCAs go wrong, let us first consider how pragmatic approaches fare when they come to account for weak binding examples:

(5) Most species have members that are small

A UCA for (5) would offer the following analysis of the sentence's semantic content relative to contexts of utterance, *c*:

(6) [Most *x*: *x* is a species] *x* has members that are small *<relative to the salient comparison class at c>*



For reasons already discussed, such an analysis of (5) seems to make adequate predictions of the truth-conditional content speakers are able to associate with utterances of (5). That is, the advocate of the pragmatic approach will effectively take it that the comparison class for ‘small’ is always fixed relative to a context of utterance. So, cases where we interpret the implicit gap associated with ‘small’ as if it were *free* (i.e. not bound) are simply cases in which the contextually salient comparison class is not the set of individuals denoted by the NP embedded in the higher quantifier phrase (i.e. ‘species’); whereas cases where we interpret the implicit gap *as if* it were bound by the higher quantifier phrase are cases where the contextually salient comparison class happens to be the very set of individuals denoted by the embedded NP.<sup>42</sup>

However, pragmatic approaches face real problems when it comes to accounting for strong binding examples. For example, consider (7):

---

<sup>42</sup> Stanley ultimately rejects the view that pragmatic approaches can account for weak binding examples. His argument rests on the claim that, in order to account for the so-called bound reading of (5), the advocate of the pragmatic approach is essentially arguing that ‘pragmatic processes can supply variables’ (Stanley 2007, p.194). This certainly fits well with Recanati’s suggestion that we might account for instances of what I have called semantic expansion by arguing that such instances involve the implementation of an *expansive variadic operator*, which effectively serves to increase the adicity of the predicate to which it is applied, thus producing an extra variable in the argument structure of that predicate. Recanati is hardly committed to this thesis though. Without going into the reasons why Stanley thinks all this is a bad thing, the problem with his claim is that it seems to be false of pragmatic approaches as he describes them. That is, it seems to me that the advocate of the pragmatic approach need not take it that the so-called *bound reading* is due to binding at all. Instead, it could simply be said that in cases such as (5), the contextually determined comparison class just happens to coincide with the class of individuals picked out by the NP (i.e. ‘species’) embedded in the quantifier phrase (i.e. ‘most species’).

(7) In every room in John's house, he keeps every bottle in the corner

According to Stanley, a UCA for (7) will predict the following interpretation of the sentence relative to contexts of utterance *c*:

(8) In [every *x*: *x* is a room in John's house], he keeps every bottle *<in the domain that is salient at c>* in the corner

Stanley argues that this is an *absurd* interpretation of (7) (Stanley 2007, p.56). Not only does (8) give the *wrong* interpretation for any utterance of (7), but it also results in a pretty striking grammaticality violation due to the fact that it involves an instance of vacuous binding. That is, 'in the case of [(7)], the quantifier 'every room' must bind a variable in the syntactic structure of the sentence 'every bottle is in the corner', on pain of ungrammaticality' (Stanley 2007, p.214). Ultimately, then, Stanley takes it that UCAs fail to provide plausible accounts of the way speakers are able to interpret utterances of a significant range of sentences. His verdict is that we should abandon approaches that appeal to UCAs, and seek alternative ways of accounting for contextual effects on utterance interpretation.

### 1.3. Logical form, the binding assumption, and aponic elements

Having shown that pragmatic approaches provide inadequate accounts of the ways speakers actually interpret utterances of various sentences, Stanley argues that we need to reject the idea that contextual information can be incorporated into interpretation via the implementation of unconstrained pragmatic processes. Accordingly, Stanley is led to argue that there must be aponic pronominal elements that are apt for binding situated in the logical forms of sentences that contain context sensitive expressions.

The logical form of a sentence, as Stanley uses the term, is its LF; i.e. a phrase marker that provides a description of the 'real structure' of that sentence (Stanley and Szabó 2000a, p.247). In postulating that there are aponic elements situated in the logical forms of sentences, then, Stanley is committing to a syntactic thesis about the actual structural arrangement of natural language sentences. As Stanley himself points out, though, 'syntactic structure cannot be postulated on semantic grounds' (Stanley 2002b,

p.368); so in order to support his thesis, Stanley needs to provide syntactic evidence for the presence of aphonic elements at LF. In order to provide such evidence, Stanley deploys *the argument from binding*.

The argument from binding is predicated on the *binding assumption* (BA):

If  $\alpha$  and  $\beta$  are within the same clause, and  $\alpha$  semantically binds  $\beta$ , then  $\alpha$  either is, or introduces, a variable-binding operator which is co-indexed with, and stands in a certain specified structural relation to, a variable which is either identical to, or is a constituent of  $\beta$ . (Stanley 2000, p.412)

What the binding assumption essentially states is that every instance of semantic binding corresponds to an instance of syntactic binding.<sup>43</sup> So, given that the binding examples are cases where some form of semantic binding occurs, binding must also occur in the underlying syntax of the relevant sentences: i.e. at LF. Consequently, if we accept the binding assumption, we must conclude that in order for semantic binding to occur in the binding examples considered, there must be elements located in the syntactic structures of the relevant sentences that are apt for binding. Since such elements can either be bound (when they occur in the domain of a binding operator) or free (when they do not), they essentially behave like pronominal elements such as ‘his’, which can be bound when it occurs in the domain of a binding operator – e.g. ‘Every man has his enemies’ – but can be free when it does not – e.g. ‘He laughed at every man’. Hence, Stanley takes the relevant aphonic elements to be *pronominal* aphonic elements.


Let us look at some of the details of Stanley’s approach in relation to an example. Stanley’s binding argument essentially shows that there is an aphonic element situated in the syntactic structure of (9a), which is then bound when that sentence is embedded in the more complex (9b):

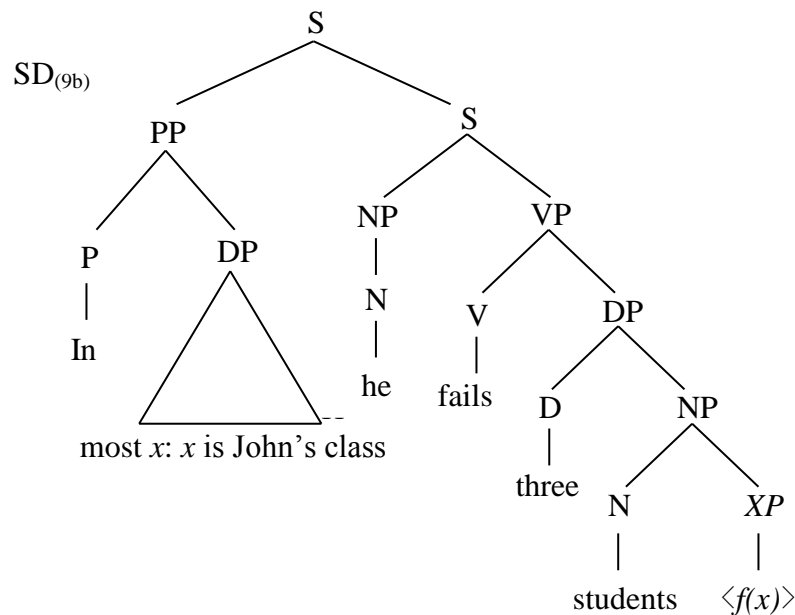
---

<sup>43</sup> For further argument to the effect that there is this close association between instances of semantic and syntactic binding, see Heim and Kratzer 1998, p.260. Note that their account varies somewhat from Stanley’s.

- (9) a. He fails three students  
 b. In most of John's classes, he fails three students

One question to ask is what position the relevant aphonic element occupies in syntax. Stanley speaks of aphonic elements as being 'associated' with context sensitive expressions. In earlier accounts, he claimed that 'a variable is associated with a syntactic element...just in case it *co-habits* a node with it' (Stanley, p.101). However, Stanley later rejects this claim. The reason for this is that if aphonic elements are taken to 'co-habit' nodes occupied by their associated lexical items, then they do not occupy independent positions in syntactic structure, but are instead conceived of as 'sub-syntactic' elements, which cannot participate in the kinds of structural relations that obtain between positions in syntactic structure, such as binding relations (see Stanley 2007, p.222, n.15). Consequently, Stanley is led to the view that aphonic elements are adjoined elements that occupy their own terminal nodes in syntactic structure. Hence, we can take the following to provide accurate representations of the LFs of (9a) and (9b):





As we see here, the aphonic element is represented as ' $f(x)$ ' (Stanley represents this as ' $f(i)$ ', but for the sake of clarity, I shall represent the aphonic element as being partly comprised of a variable element that can be bound by a quantifier phrase). In its free occurrence in (9a), the value of the aphonic variable  $x$  relative to a context is some or other contextually salient object (e.g. 'John's class') and the value of  $f$ , again relative to context, is a function from objects  $x$  to the set of objects that satisfy the contextually salient property, such as being *in* John's class. In its bound occurrence in (9b), of course,  $x$  is bound by the quantifier denoted by the DP complement of the PP 'In most of John's classes'. For the sake of clarity, I have marked the phrase formed by the aphonic element as 'XP'. That said, there is good reason to treat this as a PP adjunct, whose prepositional head is occupied by  $f$ , and whose complement NP is occupied by  $x$ .

Ultimately, then, Stanley not only provides an explanation for the kinds of semantic facts revealed by strong binding examples, but he also provides reasons for accepting the claim that such facts can be accounted for on the basis of the linguistically determined features of sentences. What is more, he provides an appealing account of how aphonic pronominal element might be represented at LF. The resulting view is that

contextual effects on truth-conditional content can be traced back to occurrences of aponic pronominal elements situated in the logical forms/LFs of sentences.

Whilst Stanley's position is thus certainly worthy of consideration, there are several concerns that nevertheless arise in relation to the feasibility of the approach he adopts. First of all, aponic elements are unlike the other implicit elements linguists tend to attribute to the syntactic structures of sentences (e.g. PRO elements, traces, null operators, and so on). We therefore need to ask whether Stanley has really provided adequate grounds for extending the class of such *empty category* items. Second, there are questions to be asked regarding how aponic elements actually come to occupy a role in syntactic structure. As we saw, Stanley argues that such elements occupy terminal nodes that are adjoined to context sensitive expressions, but this in itself raises questions that require further investigation. In the following section, I shall pursue some of these concerns in order to assess the aponic approach.

## 2. Assessing the aponic approach

My aim here is to raise several questions for Stanley's aponic approach. First, I ask whether Stanley's treatment of the binding examples is the only possible way of accounting for these cases. In response to this question, I shall develop my own alternative analysis of the examples, and attempt to show that there are certain advantages for preferring my treatment to that offered by Stanley. I shall then go on to look in more detail at Stanley's claim that aponic elements occupy adjoined terminal nodes. The question I want to ask here is whether any good explanation can be provided of why these elements accompany context sensitive expressions in syntax.

### 2.1. An alternative account

There are good reasons for wanting to develop an alternative account of Stanley's strong binding examples. Many of these reasons stem from the observation that Stanley's aponic elements are in many ways anomalous when compared to the covert items we usually encounter in syntactic theory. As Collins (2007) emphasises, *empty category* items such as PRO are posited by linguists so as to satisfy certain general requirements of syntactic theory that are quite independent of our semantic intuitions about particular

sentences. For example, the reason the sentence *The ship sank to collect the insurance* is thought to contain a PRO item is because of the *general* requirement that all sentences must have a subject; i.e. all grammatically acceptable sentences must satisfy the *extended projection principle* (EPP).<sup>44</sup> The problem is that no comparable independent syntactic requirement can be given for the presence of the aphonic elements Stanley proposes. One way to put this point is to say that Stanley's examples do not raise any well-formedness considerations. That is, the question of whether or not aphonic elements are situated at LF has no bearing on the syntactic well-formedness of sentences

---

<sup>44</sup> See Stanley (2002a, pp.152-3) and Collins (2007, pp.834-37) for two opposing treatments of this example. Ultimately, there are various kinds of syntactic considerations that motivate the postulation of PRO elements that do not seem to apply to Stanley's aphonic elements. One such consideration stems from the *theta-criterion*, which states that every lexical item must assign each of its  $\theta$ -roles to one, and only one, argument, and that every argument must be assigned one, and only one,  $\theta$ -role (see Haegeman 1994 for further discussion of the theta-criterion). Now, in the case of the sentence 'Bill decided to leave', the verb 'decide' *and* the verb 'leave' both need to assign an AGENT  $\theta$ -role to an external argument. That said, if we take it that the sentence does not contain a PRO element, it becomes very difficult to see how the sentence would satisfy the theta-criterion. One might argue that the subject NP 'John' actually occupies the external argument position of the infinitive VP at an earlier point in the syntactic derivation by which the sentence is generated, and so would have been assigned the AGENT  $\theta$ -role of 'leave' before then undergoing movement to the sentence's subject position. The problem, though, is that this would violate the theta-criterion, as 'John' would effectively be assigned more than one  $\theta$ -role (i.e. it would be assigned the AGENT  $\theta$ -role of 'leave' *and* 'decide'). So, if we want to maintain the theta-criterion, and also maintain that 'Bill decided to leave' is a grammatically acceptable sentence (in that it does not violate syntactic principles), we need to accept the suggestion that there is an implicit PRO element situated in the subject position of the embedded infinitival clause.

such as (10). So there is no obvious *syntactic* constraint that hangs on the question of whether an aphonic element is present or not; i.e. the absence of such an element would not lead to violations of any independent syntactic principles. As Pupa and Troseth write:

From a syntactic perspective, we find the binding argument peculiar. It never invokes any well-formedness considerations at all. And while the syntactic analysis of a sentence may ultimately have interpretational effects, one shouldn't resolve interpretational puzzles by positing syntactic constituents without developing a well-formedness argument in their favor. The binding argument, however, makes no such gestures. As such, it's difficult to imagine the binding argument grounding any syntactic thesis. This we note as a theoretical challenge for the binding argument. (Pupa and Troseth 2011, p.194)

As Pupa and Troseth go on to argue, the postulation of aphonic elements does not seem to be licensed by any independent syntactic principles. This is one reason for thinking that it would be preferable not to have to appeal to such elements, and so to look for alternative ways of accounting for the binding examples.

Now of course, Stanley's claim is that a sentence such as (10) would be *semantically* anomalous if we did not postulate the presence of an aphonic element at LF:

(10) In most of John's classes, he fails three students

That is, Stanley thinks that if we assume that the quantifier 'three students' is merely restricted to whatever happens to be the salient domain at any given context, the resultant reading would be 'one on which the second part of the sentence is completely unrelated to the first part of the sentence. Indeed, it is not clear, on a pragmatic approach, that sentences such as [(10)] express coherent propositions at all' (Stanley and Szabó 2000a, p.243). Ultimately, then, even if the absence of the aphonic element would not lead to any grammaticality violations, the sentence would nevertheless be uninterpretable if we were to reject the claim that there is an aphonic element situated in its LF. Now, it certainly does appear that when we come to interpret (10), we cannot help but understand the sentence in such a way that the PP must be related to – in that it



in some way modifies – the second part of the sentence. What is questionable, however, is that the relation the PP bears to the embedded sentence *must* be one on which it restricts the domain of the quantifier phrase ‘three students’.

We can conceive of a case, for instance, where John the teacher has been asked to attend evening classes designed exclusively for teachers in order to make sure that their knowledge of recent teaching methods is up to date. Finding these mandatory classes very dull, John takes with him a bunch of papers from his day job and sits marking them instead of properly participating in the class. It just so happens that John fails three of his daytime students in most of these evening classes. In this scenario the domain occupied by the three students is not restricted to the domain of students *present in most of John’s classes*. Instead, it is the act of *failing three students* that occurs in most of John’s classes a reading that can be represented somewhat simplistically as follows:

(11) In [most  $x$ :  $x$  is John’s class], [<sub>S</sub> he [<sub>VP</sub> fails three students] [<sub>PP</sub> <*in x*>]]

One might complain that it would be highly misleading to utter (10) in order to describe the above scenario, since Stanley’s interpretation is the one we would all tend to arrive at on hearing such an utterance. This, however, would depend heavily on how familiar we are with the relevant context, and besides, the purpose of this (strong) example was not to say that we are *most likely* to interpret the sentence in a way that demands the presence of a covert domain variable, but that we can *only* interpret the sentence in this way. What is more, a change of verb could perhaps make the interpretation I am offering here seem the more natural one; e.g. *In most of John’s classes, he emailed three students*. It seems to me peculiar that we would restrict the set of students emailed to the set of students present in most of John’s classes. Again, though, this is not to say that such an interpretation of this sentence is unavailable – exchanging the verb simply serves to make the other interpretation more palatable.<sup>45</sup>

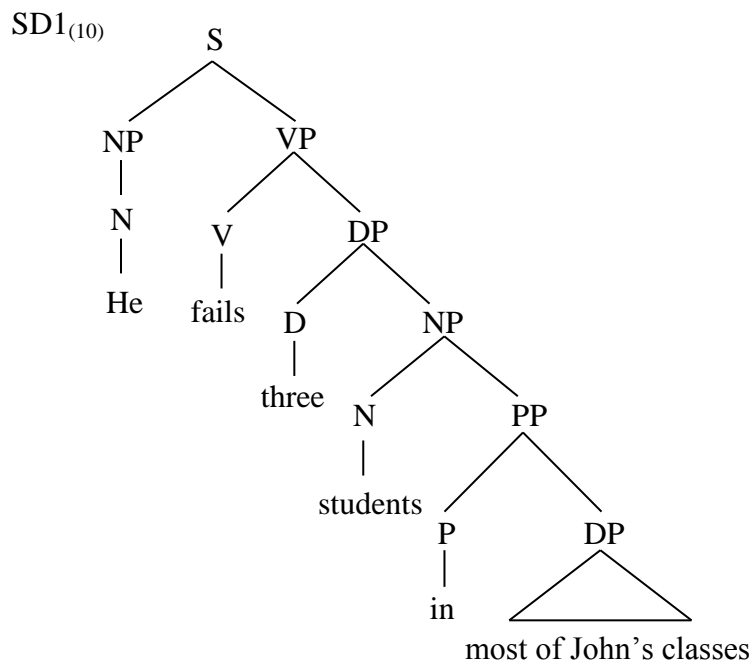
Taking the interpretation offered above to be sound, the question is how we account for the fact that the quantifier complement of the PP ‘In most of John’s classes’ seems to be

---

<sup>45</sup> This argument was, to the best of my knowledge, first set forth in Davies (2010).

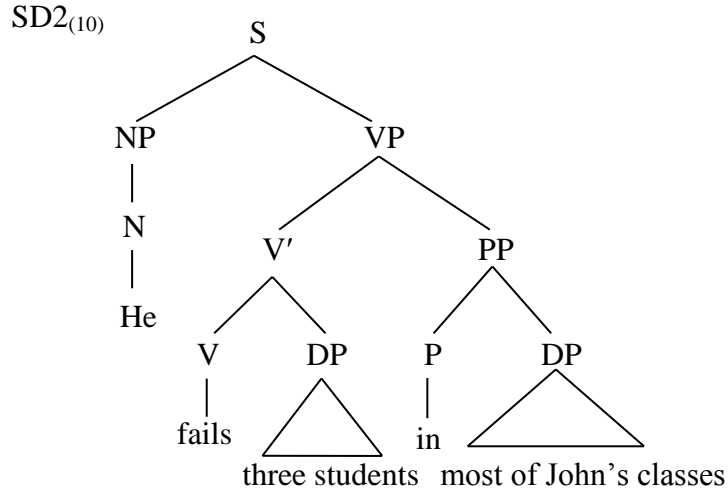
able to restrict both the location in which the act of *failing three students* took place in one instance, and the location of *the three students failed* in another. One possibility here is to posit another domain variable in the logical form of (10) whose value restricts the VP rather than the quantified NP. However, partly for the reasons mentioned above regarding the syntactically odd nature of Stanley’s variables, and partly because Stanley only ever mentions cases of domain restriction associated with quantified NPs, it may still be worth considering whether there might not be another way to account for the multiple possibilities for domain restriction here.

One option is to propose that the PP is a preposed adjunct, and that there are alternative positions it could have moved from.<sup>46</sup> The following structural descriptions depict these alternatives:




---

<sup>46</sup> Apparently Kent Bach has made this suggestion to Stanley, and whilst Stanley himself characterises it as ‘initially promising’, he immediately rejects it for reason I discuss shortly (Stanley and Szabó 2000a, p.245).



These structural descriptions do not represent two competing LF representations of (10), but rather they represent some earlier point in the sentence’s derivational history, prior to the movement of the PP. Following certain assumptions from GB theory (Chomsky 1981), we can thus treat the above representations as D-structure representations of (10). The movement of the PP occurs between D-structure and S-structure, which according to GB, is that level of representation at which the derivation ‘splits’, with further operations applying in order to attain LF and a PF representations.<sup>47</sup> If it is the case that our being competent in understanding (10) involves our being able to recognise that the sentence could have resulted from either movement, and that the PP occupied either of these positions in the underlying structure of (10), then the PP itself is what does the job of restricting either the domain of students or the location of the act of failing. This kind of account deals with both interpretations of (10) and, what is more, does so in conventional syntactic terms.<sup>48</sup>

---

<sup>47</sup> At LF, both quantifiers would both undergo a further movement to the front of the sentence leaving behind traces in order to capture the appropriate scope relations between the two.

<sup>48</sup> Of course, the notions I am appealing to here may be somewhat dated and developments in syntax have come a long way since GB theory. However, what I’m

One might wonder why the treatment of the sentence I have offered does not result in vacuous binding. I agree with Stanley that vacuous binding leads to uninterpretability, but, since the reading I have offered for (10) is perfectly coherent, I take it that no vacuous binding occurs on this reading. I think that this can be explained as follows. As has been noted, when the PP is preposed to the front of the sentence, its complement quantifier phrase is still not done moving. So, let us assume that in the transition to LF, this phrase will raise into a still higher position. Hence, simplifying somewhat, what we end up with at LF is something like the following:

SD<sub>(10)</sub> [<sub>S'</sub> Most  $x$ :  $x$  is John's class [<sub>S</sub> in  $x$  [<sub>S</sub> he fails three students]]]

Clearly, there is no instance of vacuous binding here, as we should expect given the coherence of the interpretation of the sentence I have offered.

According to Stanley, however, there is a striking problem with the approach I have adopted. As we read (10), it is unquestionable that 'John' and 'he' can be co-referential, but in SD<sub>1(10)</sub> and SD<sub>2(10)</sub>, the possibility of co-reference is ruled out. That is, we cannot understand 'John' as being co-referential with 'he' if 'John' does not occupy a position that dominates 'he' (see chapter 1 for a definition of the dominance relation). This reflects a general feature of binding, encapsulated by Condition C of binding theory, which states that an R-expression cannot be bound.<sup>49</sup> This, it seems, forces a kind of dislocation between (10), and SD<sub>1(10)</sub> or SD<sub>2(10)</sub>. According to Stanley, then, there is no way the former could have been derived from either of the latter via movement (Stanley 2000, p.420).

---

essentially trying to do is offer a treatment of examples like (10) that is not at odds with the conventional notions employed in syntactic theory.

<sup>49</sup> R-expressions, such as proper names, are *free* in the sense that they refer independently without having to be associated with any antecedent expression in the sentence in order to do so. They are distinguishable from pronominal expressions, which can be interpreted as free or bound, and anaphoric expressions, which are always bound (see Chomsky 1986, p.164).

I am not sure that this should deter us from pursuing what is otherwise a plausible approach to the above examples. Indeed, the traditional take on Condition C is that the only levels at which it *must* hold are S-structure and LF (Chomsky 1986, pp.164-86; for a classic textbook treatment of this issue, see Haegeman 1994, pp.341-45). If this is correct, there is no reason to reject SD1<sub>(10)</sub> and SD2<sub>(10)</sub> as likely D-structure representations for (10). There are, though, numerous questions about the status of the binding principles, about the level at which they should be thought to apply, about how to deal with *various* sentences that seem to violate them (e.g., *Though he was a coward, Tom fought on anyway*), and so on.<sup>50</sup> Until Stanley says more about these issues, it will remain unclear whether his appeal to binding violations rules out the treatment of his examples offered here.<sup>51</sup>

---

<sup>50</sup> The questions mentioned here multiply when one considers recent developments in syntactic theory. For instance, if the notion that indices are actually parts of syntactic objects is abandoned (see Chomsky 1995, p.228), this has deep repercussions for how to best characterise binding conditions, crossover effects, and so on. See Hornstein, Nunes and Grohmann (2005, pp.248-85) and Lasnik and Uriagereka (2005, pp.210-17) for further discussion of these problems in a contemporary light.

<sup>51</sup> In the case of the sentence ‘In every room in John’s house, he keeps every bottle in the corner’, I am again not convinced that Stanley’s construal of the sentence is the only possible interpretation. Imagine that Bill and Mary, both of whom are enthusiastic bottle collectors, are talking about John’s impressive collection of Coca-Cola bottles. In the context of their discussion, Bill might utter the above sentence, and so say something that is true iff in every room of John’s house, he keeps every bottle *ever manufactured by the Coca-Cola company* in the corner of that room. I shall not go into details of how to account for this interpretation of the sentence here, but taking it that this would be an intelligible interpretation of an utterance of the sentence, it once again appears that Stanley’s account is effectively committed to the existence of an interpretive constraint where there in fact is not one.

Now, having said all this, whether my analysis of these examples is correct or not, the point the above discussion has served to establish is that the intuitive gap associated with the quantifier phrase ‘three students’ in (10) can be saturated relative to context after all. So it no longer seems as if pragmatic approaches have any real problem accounting for strong binding examples, despite initial appearances. That said, if Stanley’s main argument against appeals to UCAs is that such appeals fail to count for strong binding examples, we are now left asking for further reasons to reject such analyses of truth-conditional content. In fact, I think that the readings I have offered for Stanley’s binding examples actually speak in favour of the idea that the kinds of gaps in content brought about by context sensitive expressions are saturated by means of the implementation of pragmatic processes.

One last thing we might consider is whether weak binding examples alone give sufficient cause to posit aphonic elements in logical form. The problem is that it is hard to see how these examples demonstrate *syntactic* grounds for doing so. We can account for the syntactic well-formedness of the sentence *Every student answered every question* with no mention of aphonic elements, and so it seems the only reason for positing such variables in this case would be to make our conception of syntactic structure square with our intuitions about the truth-conditional contents of particular utterances of the sentence. This hardly counts as giving evidence for the presence of such elements.

Of course, one could argue that weak binding examples *are* effective as long as we accept Stanley’s binding assumption. This, though, indicates that these examples only serve to demonstrate how the assumption might be applied, and it remains entirely open whether the assumption itself is to be accepted. Ultimately, the *positive* argument Stanley provides in support of accepting the binding assumption comes from demonstrating the indispensability of aphonic elements when accounting for the interpretability of certain sentences; i.e. those that give rise to strong binding examples. If this is so, and if the above argument shows that the strong binding examples are flawed, then we have reason to question Stanley’s view that logical form, conceived of as syntactic structure, is answerable to our intuitions about semantic content as he suggests.

To summarise this section, then, I have argued that we should be critical of the idea that Stanley's treatment of the strong binding examples is the correct one. In particular, I identified the syntactically anomalous character of Stanley's aponic elements as a reason to avoid postulating their presence at LF, and demonstrated an alternative way of dealing with the examples that makes no recourse to such elements but instead relies purely on the notion of movement, common to syntactic theory. If this alternative treatment works, as I think it does, Stanley needs to provide new support for his claim that the truth-conditional effects of quantifier domain restriction can be traced back to elements in logical form.

It is important to note that whereas the arguments I have produced here pertain only to instances of the strong binding examples that involve quantifier domain restriction, there are other examples Stanley discusses that are perhaps not amenable to the kind of treatment I have applied to the cases discussed here.

I now want to turn to one final of Stanley's binding examples. In doing so, I want to raise some concerns for Stanley's notion that aponic elements occupy adjoined terminal nodes in syntax.

## 2.2. 'Rain' again

The final binding example I want to consider in this section is (12):

(12) Everywhere I go it rains

Again, in the case of (12), Stanley claims that the only available interpretation of the sentence is as follows:

(13) [Every  $x$ :  $x$  is where I go] it rains  $\langle at x \rangle$

It seems to me that (12) is not a good candidate for to the alternative treatment of some of Stanley's examples that was offered in the previous subsection. That is, I think it would be at least questionable to suggest that the DP 'Everywhere I go' has moved from a lower position at which it modifies 'rain'. On the other hand, I am not convinced that

this particular binding example is as strong as Stanley thinks it is. The following example should serve to illustrate the point.

Imagine that John has been sent on a long pilgrimage by his tribe who believe that his journey will put a stop the endless downpour of rain in their village. John might travel to various holy relics, believing all the while that if he travels long and far enough, and visits enough of these relics, the gods will smile on him and end the rain in his village. Having visited the final relic, though, John receives news that it is still raining in his village. At this point, he looks up into the skies and shouts: “I give up. Everywhere I go, it rains!”. It seems to me that it is possible to get an interpretation of the truth-conditional content of John’s utterance such that what he says is true iff everywhere John goes, it rains *in the village* (i.e. at the contextually salient location). Note also, that if John shouted instead – “Everywhere I go, it rains in the village!” – we certainly do not seem to end up with an instance of vacuous binding, so the example is not strong if we take the threat of vacuous binding to be what sets apart strong binding examples. So, I think there are at least reasons to reject the claim that the only way to interpret (12) is as involving some form of binding.

Setting that point aside, then, I now want to raise questions for the claim that ‘rain’ is always accompanied by an adjoined aponic locative element that occupies a terminal node in syntax. I want to show that there are various problems for this idea, and that, ultimately, it is very hard to explain why it is that aponic elements accompany context sensitive expressions in syntax.

The basic problem is this: adjuncts are typically taken to be optional elements that are not specified in the lexical entries of expressions. That is, there are good reasons to believe that adjuncts do not form part of the entourage of lexically specified elements that must accompany lexical items in syntactic structure. Rather, adjuncts are optional elements, which can be freely added to sentences in order to modify phrases, but which are not required in order to satisfy any of the grammatically relevant requirements lexical items project onto the syntactic configurations in which they occur. Hence, once we adopt the view that aponic elements are best construed as adjuncts (as opposed to arguments), I think it becomes difficult to explain why such elements systematically



accompany context sensitive expressions in syntax. But Stanley's appeal to aphonic elements is precisely meant to account for the systematicity of extra-linguistic contextual effects on content; i.e. he wants to explain why speakers *always* take context sensitive expressions to be accompanied by certain intuitive gaps, gaps for restrictor domains in the case of quantifier phrases, gaps for locations in the case of weather verbs, gaps for comparison classes in the case of comparative adjectives, and so on. The problem, then, is that aphonic elements, *qua optional* adjuncts, do not seem to be good candidates for explaining seemingly mandatory, systematic contextual effects on interpretation.

What the above observations seem to indicate is that in order for aphonic elements to play the kind of explanatory role they are meant to play, they had best be conceived of as mandatory elements specified in the lexical entries of context sensitive expressions. That is, it seems that we need to fall back on the notion that aphonic elements are arguments (as opposed to adjuncts) that are associated with context sensitive expressions in the sentence that they are specified in the lexical entries for those expressions. However, I think the view that aphonic elements are adjuncts is unavoidable.


In order to make this point more concrete, let us begin by making the simple observation that when a locative PP appears in the explicit surface form, or PF, of a sentence such as 'It is raining', it seems to exhibit the kind of behaviour typically associated with adjuncts rather than arguments. To illustrate this point, consider the following examples:

- (14)
- a. It is raining in Paris
  - b. It is raining again/on Tuesday in Paris
  - c. Bill loves Mary
  - d. \*Bill loves again/on Tuesday Mary
  - e. Bill broke the vase
  - f. \*Bill broke again/on Tuesday the vase

What these examples show is that when we insert an adjunct, i.e. the adverbial modifier 'again', or the PP 'on Tuesday', between 'rain' and the locative PP that it combines with

in (14a), the result is the perfectly acceptable sentence (14b). That said, though, when we try to insert such items between ‘love’ and its internal argument ‘Mary’, as in (14d), or between ‘break’ and its internal argument ‘the vase’, as in (14f), the resulting sentences are clearly unacceptable.

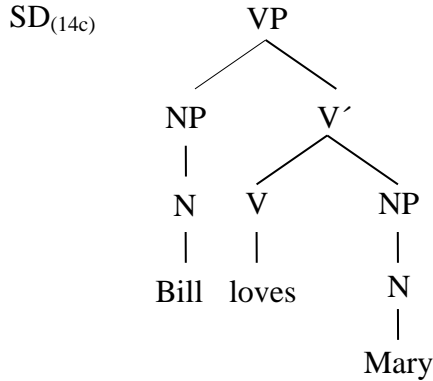
In order to understand why this happens, we need to look in more detail at the structural organisation of phrases. It is generally agreed that phrases must be configured in the following way (see Hornstein et al. 2005, pp.184-189 for discussion):



The first projection of the expression that heads the phrase is marked X, which stands for the syntactic category of the head (e.g. V for verb, N for noun, D for determiner, etc.). The intermediate projection is marked X' (or X-bar), which is that point at which the head combines with its complement (i.e. Comp). The maximal projection is marked XP, which is the point at which the head-comp configuration combines with the specifier (i.e. Spec) to form a complete phrase. A standard assumption in contemporary linguistics is that the lexically specified arguments of heads are mapped onto positions within the spec-head-comp configuration. That is, according to the conception of phrase structure adopted here, there are substantial constraints on what positions arguments can occupy. So, in the case of (14c) ('Bill loves Mary'), the verbal head maps its internal argument (e.g. 'Mary') onto its sister node, and maps its external argument (i.e. 'Bill') onto the sister node of its intermediate projection. The result is as follows:<sup>52</sup>

---

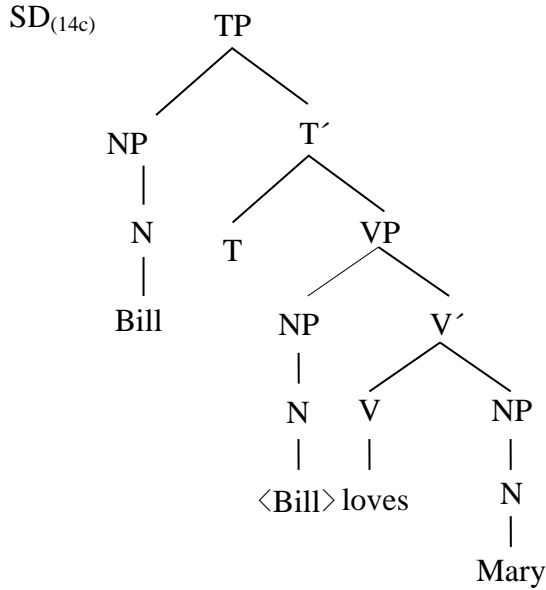
<sup>52</sup> Here, in order to simplify things, I represent 'love' as having already combined with its suffix, so as to give 'loves', within the VP.




There are two important points about this kind representation that ought to be briefly mentioned. First, according to SD<sub>(14c)</sub> (which is not a complete LF, but a partial representation that represents a certain point in the derivational history of (14c)), the subject of ‘Bill loves Mary’ is generated within the VP configuration. This is out of line with our earlier somewhat simplified representations of sentential structure. However, I make this move in accordance with the *predicate-internal subject hypothesis* (see Hornstein et al. 2005, p.81). According to this hypothesis, external arguments enter into syntactic derivations at SPEC-XP (the specifier position of the phrase). Such arguments then undergo movement to the subject position of the sentence. Second, having adopted X-bar representations of phrases, we now come to treat sentences as a certain sort of phrase – *viz.* a tense phrase (TP) – which has the same kind of structural configuration as all other phrases. So, a closer approximation to a proper LF representation of (14c) would be as follows:<sup>53</sup>

---

<sup>53</sup> The T head of TP carries inflectional features, which are realised on the V head of VP. I shall not go into the processes that ensure agreement of tense features between these elements in the syntax. Note that the triangular brackets around ‘Bill’ at SPEC-VP indicates that that item has been copied and moved to SPEC-TP (see Adger 2003 for detailed discussion on tense features and the *copy* operation).

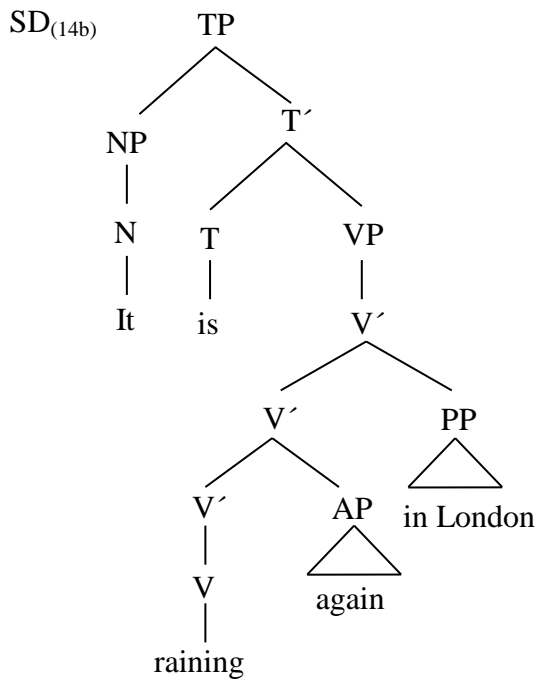


Having introduced some of the relevant linguistic notions, the important point to bear in mind is that, on this conception of the organisational principles that govern phrase structure configurations, arguments are only able to occupy very specific positions in phrasal configurations. That is, there are significant constraints on the kinds of positions arguments are able to occupy in syntax. Now, what about adjuncts? The crucial point to note about adjuncts is that they do not adjoin to phrases within the spec-head-comp configuration, but instead adjoin at a higher intermediate node. So, we can modify our representation of the structure of phrases so as to show how adjuncts adjoin to phrases as follows:



We are now in a position to state what gives rise to the kinds of interpretive effects that were listed in (14). In the case of (14c), the NP ‘Mary’, as an internal argument of

‘love’, must occupy the complement position in the VP’s phrase structure. That said, in order to attain (14d), we effectively insert an adjunct between the verbal head and its complement. As we have seen, though, only the verb’s arguments can occupy the sister node of the verbal head, so (14d) is not licensed on this conception of phrase structure. We can now ask why it is that adjuncts *can* intervene between ‘rain’ and the locative PP in (14a). Effectively, the only explanation for this is that the locative PP is an adjunct, and so can freely swap positions with adjuncts that occupy higher intermediate nodes. Bringing all this together, we can take it that the following structural description accurately represents the LF of (14b)



There are three things to note here: first, ‘rain’ does not take an external argument, so there is no SPEC-VP. The expletive ‘it’ is thus inserted directly into SPEC-TP (so as to satisfy EPP); second, since we have argued that the locative PP ‘in London’ is an adjunct, ‘rain’ does not select an internal argument, so there is no complement in the VP configuration; third, as adjuncts, the adverbial modifier ‘again’ and the locative PP ‘in London’ can be freely swapped around; i.e. it is fine to say ‘It is raining again in London’ and ‘It is raining in London again’.

It seems to me, then, that since the explicit locative PP that modifies ‘rain’ in (14a) *must* be construed as an adjunct, there is then good reason for thinking that Stanley’s locative aphonic elements would also be an adjunct. But then, to repeat the point made earlier, aphonic elements, *qua optional* adjuncts, do not seem to play the kind of explanatory role Stanley wants them to play.

With all this in mind, though, it might be argued that there are certain verbs that appear to select obligatory adjuncts. That is, there can be items that occupy adjunct positions in sentences, but which can nevertheless be taken to be specified in the lexical entries for the verbal heads with which they combine. For example, Baker (2003) argues that in (15a), the PP ‘on the table’ that directly follows the direct object ‘the plate’ is an obligatory adjunct. The same can also be said for the adverb ‘carefully’ in (15b):

- (15) a. Bill put the plate on the table
- b. Bill worded the letter carefully

Now, if the pertinent elements really are adjuncts, it should be possible freely to insert an adverbial modifier between them and the verbs that select them:

- (16) a. ?Bill put the plate again on the table
- b. Bill worded the letter again carefully

It seems to me that (16a) is awkward (hence the question mark), but it is not altogether unacceptable. (16b), on the other hand, seems absolutely fine. The question, then, is this: if these verbs can select for adjuncts, then why is it that the same cannot be said for ‘rain’? That is, perhaps Stanley could argue that his aphonic pronominals are actually lexically represented adjuncts (which would give substance to the idea that context sensitive elements are ‘associated’ with aphonic elements, even though these elements seem to be adjuncts). The problem with this kind of response, though, is that obligatory adjuncts are, not to put too fine a point on it, *obligatory*, which is just to say that they *must* be explicitly represented at PF in order for the sentences in which they occur to be grammatically acceptable. This is clear from in the following (see Baker 2003, p.318):

- (17) a. \*Bill put the plate

b. \*Bill worded the letter

To sum up, then, it seems to me that the best way of making sense of Stanley's claim that context sensitive expressions, like 'rain', are associated with aphonic elements is if it could be shown that those elements are either implicit *arguments* of some sort. All the evidence I have adduced above, however, would seem to suggest that they are aphonic elements are optional adjuncts. The problem, though, is that optional adjuncts are the last kind of thing we would expect to find listed in the lexical entries of expressions. So ultimately, it seems too hard to see why the appeal to aphonic elements would help explain why context sensitive expressions appear to be systematically associated with certain gaps in truth-conditional content.

Ultimately, having taken all this into account, all we are left with in order to justify the claim that 'rain' is *in some yet to be specified way* associated with an aphonic locative element is the *felt need* for a location that comes with occurrences of the expression in sentences. As we saw right at the outset, though, there are ways to account for this that do not appeal to linguistic constraints. That is, the above observations may lead us to re-examine Recanati's claim that we can offer purely pragmatic reasons for this *felt need*.

### 3. Summary

This chapter has provided a thorough analysis of Stanley's aphonic approach. In focusing on Stanley's argument from binding, and by questioning the force of the examples that he uses to support this argument, I have attempted to show that the argument does not suffice to support the claim that all genuine instances of context sensitivity can be traced back to aphonic elements situated in the logical forms (or LFs) of sentences. Indeed, my analysis of Stanley's strong binding examples has shown that pragmatic approaches, which appeal to the pragmatic competence of speakers in order to explain certain contextual effects on interpretation, offer a viable treatment of the relevant examples, *pace* Stanley's arguments to the contrary.

## CHAPTER 5

### THE UNARTICULATED CONSTITUENT ANALYSIS

The two main approaches considered thus far can be looked upon as attempts to retain the view that for any given sentence there must be a one-to-one correspondence between the items that constitute its structure and the elements that comprise the set the truth-conditional or propositional content it expresses on an occasion of use. As we saw, advocates of semantic minimalism argue that putative mismatches between sentential structure and content can be resolved by paring down our intuitive conception of what constitutes the proposition expressed by a sentence, whereas advocates of the aphonic approach argue that the best approach is rather to expand our conception of what constitutes sentential structure so that it squares with a more intuitive conception of content.

Either way, the aim is to ensure a good fit between sentential structure and content. Having raised a number of objections to these positions, I now want to look at an approach that attempts to show that we need not assume such a strict correspondence between sentential structure and propositional content in order to maintain the idea that there is a systematic link between sentences and the content they can be used to express on different occasions.

#### 1. Unarticulated constituents

The aim of this section is to assess *the unarticulated constituent analysis* (UCA), originally set forth by Perry (e.g. 1986, 1998). My aim is to chart the development of UCA from its earlier presentation in the work of Perry, to its more recent presentation in the work of Taylor (e.g. 2001, 2007). I shall take it that both versions of UCA are committed to the following claim: the meaning of any context sensitive expression fixes what type of contextual information must be taken into account in order to determine the truth-conditional content of any sentence in which that expression occurs. What I am



interested in clarifying is how the meanings of context sensitive expressions serve this function.

In order to answer this, I begin by considering Perry's account of the way in which contextual effects on content are fixed by the meanings of context sensitive expressions. I shall show that Perry's account is ultimately based on an appeal to metaphysical considerations. I then consider some problems for Perry's reliance on metaphysical consideration that have been raised by Taylor (2007), before finishing the section by outlining Taylor's alternative version of UCA, which effectively rejects the metaphysical considerations Perry's approach relies on and focuses instead on purely lexical considerations.

### 1.1. Perry's proposal

Perry (e.g. 1986, 1998) has argued that we can explain the systematic correspondence between sentences and the propositions they are apt to express on different occasions even if we forgo the assumption that all propositional elements have correlates in sentential structure. On this view, the propositions expressed by uses of certain sentences in part comprise *unarticulated constituents*; i.e. proper elements of propositional content that do not have any correlates at *any* level of linguistic representation, but whose incorporation into content exhibits systematicity across contexts. Standard cases include the following:

- (1) a. It's raining  
b. It's raining <*in London*>
- (2) a. It's 8 a.m.  
b. It's 8 a.m. <*GMT*>
- (3) a. The mayor is drunk  
b. The mayor <*of Reykjavik*> is drunk
- (4) a. Every teacher went on strike  
b. Every teacher <*in the UK*> went on strike

When the above (a)-sentences are uttered, a speaker's interlocutors automatically recognise that interpretation calls for extra information from context; they know the type of information required and how it is to be integrated into propositional content; and they are (in general) familiar enough with the context to retrieve the very information the speaker intended to convey.<sup>54</sup> According to UCA, none of this requires that the implicit information be traceable to discrete constituents located in the syntax of the sentences uttered (e.g. aphonic elements at LF). Instead, the UCA theorist maintains that, under the right conditions, speakers can rely on their interlocutors to identify certain elements of the propositions they wish to express without their having to 'belabour those elements in syntax', or to *articulate* them at some level of linguistic representation (Crimmins and Perry 1989, p.700). In arguing that the constituents of sentences and the elements of the propositions that they are used to express need not stand in a one-to-one correspondence to each other, the UCA theorist rejects what Perry (1986) calls *the principle of homomorphic representation* (PHR), or what Crimmins (1992) calls *full articulation*.<sup>55</sup>

---

<sup>54</sup> I say 'in general' because certain factors can militate against the retrieval of the apposite information; mistakes can occur, and speakers are sometimes required to make explicit what they initially left implicit. The same margin for error is permitted whether the relevant elements in propositional content are construed as the context relative denotations of aphonic indexicals or genuine unarticulated constituents in Perry's sense of the term. That is, even if there is (say) an aphonic indexical for locations associated with the verb 'rain', one can still make a mistake in figuring out *which* location the speaker had in mind when he or she uttered the sentence 'It is raining'.

<sup>55</sup> PHR is associated with a strict version of the compositionality requirement, according to which the semantic content of a sentence is a syntactic function of the contents of its constituents. Unarticulated constituents are not denoted by sentential constituents, and so their presence in propositional content violates this traditional notion of compositionality.

The question to ask is how unarticulated constituents come to occupy a stable role in utterance interpretation. More specifically, we need to ask how it is that when speakers encounter utterances of certain sentences, they automatically recognise the need to incorporate specific pieces of information from context into their interpretation of propositional content even though none of the constituents of the sentences uttered stands for those pieces of information. Put this way, it appears as if the UCA theorist needs to explain why a very particular set of interpretive effects routinely accompany uses of sentences despite not being controlled by the formal, context-independent features of those sentences.

Advocates of UCA, however, do not deny that the introduction of unarticulated constituents into propositional content is triggered in part by the stable linguistic features of sentences. That is, whilst the UCA theorist denies that unarticulated constituents correspond to discrete sentential constituents, he does not thereby commit to the view that the identification of such elements is due solely to unconstrained, or *free*, pragmatic processes. Rather, it is argued that unarticulated constituents are incorporated into content via *content-supplemental* uses of context (Perry 1998, 2001), which operate under constraints imposed by the interaction of linguistic *and* contextual factors:<sup>56</sup>

[T]he unarticulated constituent is not designated by any part of the statement, but it is identified by the statement as a whole. The statement is *about* the unarticulated constituent, as well as the articulated ones. (Perry 1986, p.209)

---

<sup>56</sup> Perry (2001, pp.45-46) suggests a more accurate phrase to describe what she refers to as *post-semantic uses of context* would be *post-meaning but pre-content uses of context*. To avoid the misleading connotations of the former term and the awkwardness of the latter, I'll use the phrase *content-supplemental uses of context*, or simply *content supplemental processes* (CSPs).

It would be misleading...to say that, in the case of unarticulated constituents, the context alone does the job. The whole utterance – the context and the words uttered – is relevant to identifying the unarticulated constituent. Thus, a change in wording can affect the unarticulated constituent, even though it is not a change in an expression that designates that constituent...The important principle to be learned is that a change in wording can precipitate a change in propositional constituents, even when the words do not stand for the constituents. (Crimmins and Perry 1989, p.700)

The idea, then, is that the identification of unarticulated constituents is driven in part by features of the sentences we use and in part by features of the contexts we use them in. So, to reiterate the initial question, we need to ask how utterances of particular sentences *call for* propositional elements none of their constituents *stand for* (Crimmins and Perry 1989, p.701).

Perry's (1986) proposal is as follows. The kind of information stored in the lexical entries for simple expressions is split into two distinct components; *viz.* the syntactic and the semantic. The syntactic component specifies the structurally relevant properties of expressions, such as which *argument positions* (if any) they project into syntactic structure. These argument positions must be filled by items present in syntax to satisfy grammaticality requirements. The semantic component specifies semantic rules for assigning denotations to expressions relative to contexts, where potential denotation types include the usual range of individuals, times, locations, and *n*-place relations. Expressions that denote relations are associated with a fixed number of *argument roles*, which are determined on the basis of the arities of the relations denoted.<sup>57</sup> The argument roles associated with the expressions that form a given sentence must either be saturated by elements denoted by other expressions in the sentence, or by elements identified from context, in order to attain a complete propositional content. Hence,

---

<sup>57</sup> I shall say that an expression is *associated* with a given set of argument roles iff it denotes a relation that takes the relevant set of arguments.

expressions that denote relations have both a *lexical adicity*, which specifies how many syntactic argument positions the expression projects, and a *conceptual adicity*, which specifies how many argument roles are associated the expression.<sup>58</sup>

The reason one must sometimes look to context in order to saturate argument roles is that the lexical and conceptual adicities associated with relational expressions (e.g. verbs, adjectives, and so on) are not always equal. Hence, in some cases, composition of the denotations of the expressions that form a grammatically acceptable declarative sentence will not yield a complete proposition, since filling all argument positions projected by the expressions present in a sentence does not ensure the saturation of all the argument roles associated with its constituent expressions. When this occurs, we have to utilise our knowledge of context in order to fulfil the requirement that argument roles must undergo saturation, and so attain a complete proposition:

On the way I look at things, relations have *argument roles* or parameters. These are to be distinguished from the *argument places* or *variables* that predicates that express the relations may have. My picture of unarticulated constituents is that there are argument roles that are not represented by explicit argument places. We fill the *argument role* which is filled from context. (Perry 2001, p.47)

This leaves us with the following view of utterance interpretation: the meanings of expressions are encoded in lexical rules whose application (relative to contexts) results in the assignment of denotations to uses of those items. When the *semantic phase of*

---

<sup>58</sup> The argument positions lexical items induce in syntactic structure constitute their *argument structure*. A lexical item's argument structure specifies the arguments it takes by syntactic category. Perry's argument roles are presumably specified in a manner that makes explicit what types of propositional elements can saturate the relevant roles in propositional content. Hence, the verb 'hit' encodes two argument positions for DPs (assuming that argument structure encodes external arguments), and denotes a relation with two argument roles for individuals.

*interpretation* (Perry 1989), which involves implementing lexical rules so as to assign semantic values to expressions, results in the assignment of an  $n$ -place relation to some constituent of the sentence uttered, but none of the sentence's other constituents are assigned denotations fit for saturating one or more of the argument roles of that relation, content remains incomplete. Interpreters must then proceed to the *content-supplemental phase of interpretation*, utilising their knowledge of context to saturate empty argument roles and so attain a complete content.

The type and number of propositional elements identified from context at this stage is determined by the type and number of argument roles that stand in need of saturation. On some occasions, a speaker's interlocutors might identify the wrong element from context, and so misinterpret the speaker's utterance.<sup>59</sup> On others, they may fail to identify *any* element of the appropriate kind, and so fail to interpret the utterance altogether. When the correct element is identified from context, however, the proposition that the speaker intended to convey is determined, and successful communication achieved. It is in this way that the interaction of the words we use and our ability to derive information from the contexts in which we use them brings about the systematic incorporation of unarticulated constituents into propositional content. The relevant interpretive phenomena are thus accounted for without assuming the strict

---

<sup>59</sup> When one identifies the wrong element from context, one fails to identify an unarticulated constituent of the proposition the speaker intended to express. A proposition is still determined, however, so perhaps one might still speak of the element identified as an unarticulated constituent, albeit one of a proposition the speaker did not intend to express. Perhaps Perry would find such talk infelicitous, but we can let that pass. The point is that the conditions that must be in place in order for an unarticulated constituent to be identified during interpretation, i.e. the conditions that speakers and interpreters alike exploit in order to communicate in lieu of full articulation, allow for some degree of slippage between the proposition one intended to express and the proposition grasped by one's interlocutors.

compositionality requirement associated with PHR (Neale 2007; Crimmins and Perry 1989).

## 1.2. The lexical and conceptual adicity of ‘rain’

Perry’s UCA approach has been taken to provide a particularly convincing explanation of what gives rise to the requirement that the interpretation of weather reports involves the identification of locations from contexts. Consider (5):

(5) It is raining

On Perry’s (1986, 2001) analysis of (5): ‘it’ is a non-referring expression required for purely grammatical reasons; ‘is’ refers to the time of utterance; and ‘rain’, which projects one argument position for tense, denotes a dyadic relation between times and locations – it being the nature of raining events that they occur at particular times at particular places. Syntactically speaking, combining these expressions as in (5) yields a well-formed sentence, yet, semantically, composition of the denotations of the relevant expressions fails to yield a complete proposition, since the *rain* relation requires an argument that is not supplied (denoted) by any of the constituents of which the sentence is comprised.

Put another way, the predicate ‘rain’ has just one syntactic argument position, filled here by the appropriate tense morphemes (‘is’/‘-ing’), but denotes a relation with two argument roles, one of which is here left empty. So, although ‘we have the syntax of the sentence “It is raining” and the meanings of each of the component words, we still don’t have the content’ (Perry 2001, p. 45). A location must thus be identified from context to saturate the empty argument role and attain a complete proposition (note: this need not be the utterance location (Perry 1986, pp.210-211)).

Now, there is no doubt that such an analysis provides an intelligible explanation of why it is that the interpretation of weather reports calls for the identification of locations from contexts, and thus serves to demonstrate ‘the theoretical possibility and coherence of truly unarticulated constituents’ (Perry 1998, fn.4). The question, though, is whether there are good reasons to accept the claim, upon which the analysis is predicated, that

there is a discrepancy between the number of argument positions enumerated in the lexical adicity of ‘rain’ and the number of argument roles enumerated in its conceptual adicity. In order to assess this, we need to examine whether the analysis offered provides a plausible characterisation of the lexical and conceptual adicities of ‘rain’.

Focussing initially on the syntactic aspect of Perry’s analysis, he is quite right to say that the expletive subject of the sentence is not a syntactic argument of the predicate. Indeed, expletives do not fill argument positions projected by predicates, but are required to satisfy the structural requirement that clauses must have subjects; i.e. the extended projection principle (EPP) (Chomsky 1981).<sup>60</sup> As we saw in the previous chapter, expletives are thus inserted into syntactic structure outside of VP configurations within which the syntactically relevant argument taking requirements of verbs are resolved. To this extent, Perry’s analysis gets the lexical adicity of ‘rain’ right.

Having said that, though, current thinking in linguistics militates against Perry’s further claim that ‘rain’ encodes an argument position for tense. It has already been shown that, according to the consensus view on such matters, what we normally think of as sentences are actually tense phrases (TPs), whose (functional) heads take VP complements. So, again, on the assumption that verbs combine with all of their arguments *within* the VP configuration, tense heads (which carry inflectional morphology) combine with verbs only *after* the argument taking requirements of the latter have already been satisfied (i.e. after the verb’s argument positions have been

---

<sup>60</sup> Neale says that in the case of (6), there is ‘exactly one argument role occupant ([e.g.] Reykjavik) with no corresponding argument position occupant, and exactly one argument position occupant (‘it’) with no corresponding argument role occupant’ (Neale 2007, p.322). This is fine as long as we bear in mind that the argument position mentioned here is not one listed in the lexical adicity of ‘rain’.



filled).<sup>61</sup> This is one way of showing that the lexical entry for ‘rain’ does not list an argument position for tense, which means that we can take ‘rain’ to have a lexical adicity of zero. I take it that although this indicates that there is a greater discrepancy between the lexical and conceptual adicities of ‘rain’ than Perry suggests, it does nothing to detract from the remainder of his analysis.

Moving on to the semantic aspect of Perry’s analysis, we need to consider what grounds there are for the idea that ‘rain’ has the conceptual adicity attributed to it. This leads to more general questions regarding what determines the conceptual adicities of lexical items. Intuitively, one might well assume that knowledge as to which argument roles are associated with which expressions stems from knowledge of the lexicon; i.e. that the conceptual adicities of expressions are lexically specified. However, according to the UCA approach, although lexical entries specify *which* relations expressions denote, they specify no information pertaining to the adicities of those relations. Instead, such things are determined on the basis of extra-linguistic, essentially metaphysical, considerations pertaining to the way the world must be in order for a given relation to obtain. So, on this view, conceptual adicity is determined on the basis of metaphysical considerations. In the case of ‘rain’, then, although the expression’s lexical entry specifies a specific relation as its denotation, knowledge of the adicity of that relation does not derive from knowledge of the lexical item *per se*, but from knowledge of the fact that raining events must occur at particular times at particular places.<sup>62</sup> If this is right, our knowledge of

---

<sup>61</sup> Recall that according to the VP internal subject hypothesis (PISH; see Koopman and Sportiche 1991), the subjects of verbs are actually taken to combine with their verbs within the VP configuration, at SPEC-VP, before *moving* to SPEC-TP. In order to preserve PISH in the face of complications regarding ditransitive constructions, it has been proposed that VP configurations are dominated by the functional projections of light verb (*v*) heads (see Larson 1988, Chomsky 1995, Hornstein 2005).

<sup>62</sup> Perry is quite explicit that his approach relies on such a distinction between the role grammatical and metaphysical considerations play in interpretation. For example, in

which argument roles are listed in the conceptual adicities of expressions stems from aspects of our general world knowledge that are not encapsulated in the lexicon.

### 1.3. Rejecting metaphysical considerations

The notion that the number of argument roles associated with a given lexical item is established on the basis of metaphysical considerations is generally rejected by contemporary attempts to defend the UCA approach (e.g. Taylor 2001; Neale 2007; Korta and Perry 2011). A chief problem for Perry's earlier account is set out by Taylor (2001, 2007).

According to Taylor, in the case of weather reports, it does indeed seem to be the case that in order to attain a truth evaluable content, one needs to identify a location from the relevant context. Consequently, it seems safe to say that that the propositions expressed by such utterances partly unarticulated constituents. However, if the only reason for thinking this is that raining events must occur at particular locations, then it ought to follow that interpreting utterances of the following sentences also requires the identification of specific locations from contexts, since events like dancing, eating, and giving must also take place at locations:

- (6) a. Mary danced
- b. Bill ate five apples
- c. Mary gave the book to Bill

The problem for Perry's account is that there is simply no reason to think that interpreting utterances of the sentences of (6) involves identifying locations from

---

discussing the difference between the verbs 'beat' and 'win' he points out the following: 'The metaphysics of winning requires a victor and a vanquished, but the grammar of 'won' does not require we identify the vanquished' (Perry 2001, p.48). That is to say that 'won' has just one argument position for a victor, but has argument roles for both the victor and the vanquished, unlike 'beat', which has two argument positions and denotes a relation with two argument roles.

contexts. As Taylor quite rightly points out, ‘[o]ne can say something fully determinate, something fully truth evaluable, by uttering [(6a)] even if context provides no place as the place where the dancing took place’ (Taylor 2001, p.54).

That said, since one can determine the proposition expressed by ‘Mary danced’ without needing to know where she did her dancing, the proposition expressed does not contain an unarticulated locative element. So, it seems we must conclude that the metaphysical fact that dancing, like raining, must happen at a place does not ensure that there is an argument role for locations present in the conceptual adicity of ‘dance’. Put more generally, the arguments of a given relation need not all be listed in the conceptual adicity of the lexical item we use to denote that relation.

What is it, then, that makes ‘rain’ behave differently from ‘dance’ in the relevant respect? Taylor’s answer is as follows:

On the view that I favor, the verb ‘to rain’ has a lexically specified argument place that is theta-marked THEME that takes places as values. My claim is that, in the lexicon, rainings are explicitly marked as a kind of change that places undergo. But from the point of view of sentence-level syntax, such lexically specified parameters are what I call subconstituents rather than constituents. No constituent of the sentence (3) need serve as an argument place for the verb ‘rain’. Yet, despite the fact that this lexically specified argument place need not be expressed as a sentence-level constituent, it makes its presence felt by “demanding,” on pain of semantic incompleteness, to be assigned a contextually supplied value. (Taylor 2007, p.225)<sup>63</sup>

Taylor, following Recanati (1993), dubs the position outlined *parametric minimalism* (PM). Inasmuch as the approach has been widely endorsed by advocates of UCA, it requires further attention. In particular, the following three claims, which are all

---

<sup>63</sup> As Neale (2007, p.321) points out, Taylor’s talk of *argument places* here is probably best avoided, since Perry often uses the term argument place to talk of argument positions.

endorsed by Taylor, need to be assessed: (i) that *subsyntactic parameters* (i.e. argument roles that do not correspond to syntactic argument positions) are lexically specified and assigned  $\theta$ -roles; (ii) that argument roles are subsumed in lexical entries in virtue of their inherent thematic relevance; and (iii) that argument roles can be syntactically inert despite being lexically specified. I now want to explore these claims in further detail.

#### 1.4. Taylor's proposal

In order to assess PM as a modified version of the UCA approach, I start by focussing on the claim that the conceptual adicities of expressions are specified as part of their lexical entries. This marks an important adjustment to Perry's initial proposal that argument roles are associated with lexical items on the basis of metaphysical considerations. As Taylor argues:

Perry's mistake, I think, is to suppose that the bare mismatch between predicate and relation itself somehow explains why some argument roles must be contextually supplied, on pain of semantic incompleteness, while others need not be. My alternative hypothesis is that it is typically the lexical structure of the verb itself, and not language independent facts about the relation or even facts about the mismatch between predicate and relation, that determine[s] what must be supplied and what need not be supplied by context. The verb itself directs its own semantic completion by, as it were, demanding that occupants of certain argument roles, but not others, be supplied in context, sometimes, perhaps, as the value of sentence-level constituents, but not always. (Taylor 2007, p.228)

According to PM, then, knowledge of which argument roles are associated with which context sensitive expressions stems entirely from our knowledge of the lexicon. Hence, the semantic component of a lexical entry does not only indicate which relation an expression denotes (as per Perry's earlier view), but also lists a set of argument roles that are pertinent for the interpretation of the relevant expression, and thus must undergo saturation during interpretation 'on pain of semantic incompleteness'.

Now, an interesting aspect of Taylor's approach is that the argument roles enumerated in an expression's conceptual adicity are taken to undergo theta-assignment. Indeed, as

Taylor proceeds to elaborate PM, he comes to characterise the conceptual adicity of ‘rain’ in terms of the verb’s ‘lexically specified thematic structure’:

The verb [‘rain’] itself – in particular, its lexically specified thematic structure – is the source of the felt need for the contextual provision of a place or range of places where a raining happens. (Taylor 2007, p.226)

On this view, the component of a lexical entry that specifies argument roles, and the component that specifies  $\theta$ -roles (i.e. its *theta-grid*), come to be seen as one and the same. Just as some argument roles – i.e. subsyntactic parameters – do not correspond to argument positions, we now adopt the view that some  $\theta$ -roles do not correspond to argument positions. It thus follows that we can equate subsyntactic parameters with unassigned  $\theta$ -roles, and, consequently, that we can equate the requirement that subsyntactic parameters must undergo saturation (or be assigned values from context) with the requirement that expressions must assign all their  $\theta$ -roles. By equating argument roles with  $\theta$ -roles in this way, we retain a certain economy to our characterisation of the kind of information stored in lexical entries, since, given that  $\theta$ -roles must be lexically specified anyway and argument roles are assigned  $\theta$ -roles, there seems little use in duplicating the relevant aspect of lexical content.

These observations lead us to the second claim behind PM; *viz.* that argument roles are subsumed into lexical content in virtue of their exhibiting a certain thematic relevance. More fully, Taylor (2007) argues that idiosyncratic differences between the kinds of argument roles that are associated with expressions can be accounted for in terms of differences in the way the relations denoted by those expressions are lexicalised (see also Korta and Perry 2011). Lexicalisation involves a more or less *fulsome* incorporation of relations into lexical content. So, while the relations denoted by ‘rain’ and ‘dance’ both (metaphysically speaking) occur relative to locations, the former undergoes a more fulsome lexicalisation in that its locative aspect is subsumed in lexical content as an argument role.

Taylor’s idea seems to be that the reason *why* the rain relation undergoes a more fulsome lexicalisation in the relevant sense is that we conceptualise raining events as

events that *happen to* locations; i.e. we think of locations as *undergoing* raining events (which is not the case for dancing events).<sup>64</sup> Hence, since the notion of that which *undergoes an event* has long been taken to define a category of  $\theta$ -role (i.e. THEME), the locations of raining events are apt for inclusion in thematic structure. If this is on the right track, it further supports, and perhaps explains to some extent, the close link between notions of conceptual adicity and thematic structure. Argument roles enumerated in the conceptual adicities of expressions still have their source in the metaphysical features of the relations those expressions denote, but such features only make their way into lexical content in virtue of the fact that they participate in the relation denoted by the expression in a thematically relevant way.

The view we have come to is this: a lexical entry for an expression specifies a number of argument positions, and a number of  $\theta$ -roles. Some  $\theta$ -roles are assigned to argument positions. When a  $\theta$ -role is not assigned to any argument position, it makes its presence felt by demanding that a certain piece of contextual information – i.e. an unarticulated constituent – be incorporated into interpretation. Once unarticulated constituents are identified for all unassigned  $\theta$ -roles, interpretation is complete. Those  $\theta$ -roles that do not correspond to argument positions specified in the lexical adicity of a given expression are taken to be ‘locked’ in the subsyntactic basement.

Ultimately, I take Taylor’s approach to be more appealing than Perry’s, largely because it abandons the unattractive notion that appeals to metaphysical considerations can play an explanatory role in accounting for the kinds of factors that give rise to constraints on interpretation. That said, I think there are certain problems for Taylor’s proposal, to which I now turn.

---

<sup>64</sup> I take it that, on this view, the conceptual adicity of ‘dance’ does not encode a location parameter just because our understanding of the way dancings relate to their locations cannot be classified under any common category of  $\theta$ -role. Of course, if one counts LOCATION as a category of  $\theta$ -role, this explanation runs aground. However, as we shall see, there may be good reasons for limiting the range of  $\theta$ -roles specified in lexical entries.

## 2. Objections to parametric minimalism

I want to offer just one main objection to Taylor's approach. Though the objection can be stated simply, I think it severely undermines the explanatory value of PM. Basically, I do not think that the appeal to 'understood participants' helps to account for the very cases Taylor is aiming to explain.

### 2.1. Understood participants

Does the claim that locations participate in raining events, in the sense that they *undergo*, and so are the understood THEMES of, such events, really serve to explain why speakers take it that they must identify a particular location from context to interpret utterances of sentences containing 'rain'? Taylor thinks that it does, but I think there are reasons to question this.

We saw that Perry argued that the reason speakers incorporate locations into their interpretation of utterances of sentences containing 'rain' is that raining events must happen *at* a location. Taylor then pointed out that dancing events also happen at locations, but this does not lead speakers to incorporate locations when they come to interpreting utterances of sentences containing 'dance'. Taylor then argues that the difference between 'rain' and 'dance' is that whilst dancings simply happen *at* locations, rainings happen *to* locations. That is, locations are understood participants of raining events. This then leads Taylor to argue that there is a location specified in the conceptual adicity of 'rain'. What we want to know is whether Taylor's appeal to understood participants is really an improvement on Perry's appeal to metaphysical considerations.

Consider the following cases:

- (7) a. The key opened the door
- b. Bill bought flowers for Mary
- c. Mary ate

The verbs in each of these sentences stand for events or activities that are understood to involve participants that are not mentioned in the sentence. That is, in the case of (7a), we know that someone *used* the key to open the door. In the case of (7b), we know that

Bill bought flowers *from* someone, *for* some amount (otherwise, he simply did not *buy* flowers for Mary). And in the case of (7c), we know that there is some food that Mary is eating. What is more, all of these understood (but ‘off-stage’) participants participate in the event in some thematically relevant way. So for example, the person who used the key would be the understood AGENT who used the key as an INSTRUMENT to open the door, which is the THEME in the event. So, on Taylor’s account, there will be  $\theta$ -roles for these understood participants specified in the thematic structures of the relevant verbs.

Perhaps (7a) is the weakest of the cases here, as one can conceive of cases where we do not need a further AGENT participant. For instance, it may be the case that we can construe the key as the AGENT of the event, which militates against the necessity for a missing participant. I think that this is absolutely right. The key in question may be an enchanted key that acts under its own volition. That said, though, the same response cannot be given in relation to (7b) or (7c). That is, we simply cannot conceive of a buying event that does not involve a buyer, seller, amount paid, and thing bought. All of the participants of ‘buy’ are thus conceptually necessary. Similarly, if Mary ate, then there *has* to be something she ate.

Setting the quibble aside, then, the problem here is that when we come to interpret the truth-conditional content expressed by utterances of the sentences of (7), we do not feel the need to identify the individual, say, from whom Bill bought the flowers. Nor do we feel the need to identify who used the key to open the door. Nor do we feel the need to find out what it is Mary is eating. At best, we understand that there must have been *someone* who sold the flowers to Bill, *someone* who used the key to open the door, and *some food* Mary ate.

This signals an important distinction raised by Fillmore (1986). He writes:

It has occasionally been pointed out that a distinction is needed between what I will call *indefinite null complements* (INC) and *definite null complements* (DNC). With *definite null complements* the missing element must be retrieved from something *given* in the context; with *indefinite null complements* the



referent's identity is unknown or a matter of indifference. One test for the INC/DNC distinction has to do with determining whether it would sound odd for a speaker to admit ignorance of the identity of the referent of the missing phrase. It's not odd to say things like, "He was eating; I wonder what he was eating"; but it is odd to say things like, "They found out; I wonder what they found out." The missing object of the surface-intransitive verb EAT is indefinite; the missing object of the surface-intransitive verb FIND OUT is definite. The point is that one does not wonder about what one already knows. (Fillmore 1986, p.96)

In line with this, we can take it that we get *indefinite*, as opposed to *definite* readings, of the *missing elements* in the sentences of (7). What Taylor's approach is essentially meant to account for, though, is why we get a *definite*, as opposed to an *indefinite*, reading of the missing element associated with the 'rain' predicate. The problem is that Taylor's appeal to *understood participation* as the key explanatory notion here does not serve to explain why we get a definite rather than an indefinite reading.

What is more, the above observations suggest that we can question whether our ability to recognise that there are certain indefinitely construed event participants is really an ability that arises from our knowledge of the lexicon. To clarify, we know that 'Bill bought Mary flowers' is true iff Bill bought Mary flowers. We also know that Bill *bought* Mary flowers iff he paid someone (or perhaps *something*, say if he used a vending machine) a certain amount for those flowers. But does our knowledge of this latter fact really stem from our knowledge of the lexicon? It seems to me that it is at least possible that we can think of such knowledge as arising from our general conceptual competence, or world knowledge, rather than from some aspect of our linguistic competence. Hence, even though we *can* understand the indefinite participants as participating in the buying event in a thematically relevant way, this does not guarantee that the lexically specified thematic structure of the verb encodes  $\theta$ -roles for those participants. This raises questions for Taylor's assumption that lexical entries for expressions encode information pertaining to the understood participants in the events those expressions encode.

Once we begin to question the connection between that which is lexically encoded, say in thematic structure, and that which is encoded in conceptual adicity, we then have to question whether our knowledge of the fact that *raining events* appear to have locations as participants is something that stems from information encoded in the meaning of ‘rain’, or from information encoded in the concept RAIN. It seems to me that the only real reason Taylor has for arguing that such information is *lexically* represented in some way is that we have that *felt need* to identify locations from contexts when it comes to interpreting the truth-conditional content of utterances of sentences in which the ‘rain’ predicate occurs. But again, this was the very fact Taylor’s proposal was supposed to explain.

### 3. Summary

In this chapter, I have given an extensive overview of the development of UCA, from its origins in Perry, to its more recent manifestation in the form of Taylor’s PM. I have shown that advocates of UCA aim to give an explanation of certain systematic contextual effects on utterance interpretation by tracing the relevant effects back to certain aspects of the meanings of expressions that are subsyntactic, in the sense that they have no structural signature (unlike the arguments encoded in an expression’s lexical adicity). I showed that the main line of continuity between these approaches lies in the idea that the lexical and conceptual adicities of expressions need not be equal, but that both impose constraints on interpretation.

The argumentative component of the chapter has proceeded as follows. First, I outlined and agreed with Taylor’s arguments against Perry’s claim that we can determine an expression’s conceptual adicity on the basis of purely metaphysical considerations. I then showed that, on Taylor’s alternative outlook, we can determine an expression’s conceptual adicity on the basis of an appeal to the notion of understood participation. I then showed that Taylor’s appeal to the notion of understood participation does not succeed as an explanation of the kinds of interpretive effects his modified version of UCA, i.e. PM, was meant to explain. Ultimately, I take it that Taylor’s approach does not provide an answer to the question of why it is that when certain expressions appear

in a sentence, speakers have to look to context in order to determine the truth-conditional content of utterances of that sentence.

## CONCLUSION

The approaches assessed in the last three chapters of this thesis all share the assumption that the truth-conditional contents of declarative sentences are compositionally determined by their lexico-syntactic properties. That is, sentences have truth-conditions, or express propositions, relative to contexts purely in virtue of the context-invariant meanings of their constituent lexical items and their syntax. Each approach thus adopts the standard view that the role of a semantic theory is to generate pairings of sentences with their content by providing: (i) context-independent specifications of the meanings of lexical items in terms of the contributions those items make (relative to contexts) to the contents of all sentences in which they occur; and (ii) context-independent specifications of the semantic effects brought about by combining lexical items in different syntactic configurations.

Accordingly, if the content of a sentence is apt to vary on different occasions of use, the relevant variation must be reflected in our specification of its lexico-syntactic properties. Of course, there are restrictions on how much context-dependence can be specified in advance. So, as we saw, each approach aims to set a principled limit on the extent to which we need to attribute context-dependence to the lexico-syntactic properties of sentences in order to account for the contribution such properties make in determining truth-conditional content. More fully, semantic minimalists argue that we need only attribute context-dependence to standard indexicals in order to achieve adequate pairings of sentences with their content; aphonic theorists argue that we need only attribute context-dependence to standard indexicals and aphonic pronominals (at LF); and UCA theorists argue that we need only attribute context-dependence to standard indexicals and expressions whose meanings encode subsyntactic parameters.

Now whilst the approaches discussed clearly succeed in attaining some form of pairing between sentences and their content, each ends up with a different notion of what gets paired. For instance, the semantic minimalist ends up adopting a distinctly ‘thin’ conception of what constitutes truth-conditional content, whilst the aphonic theorist and

the UCA theorist both make substantial revisions to prior conceptions of what constitutes lexico-syntactic structure. I have argued that in order to assess the adequacy of the pairings achieved by each of these approaches, we have to ask whether there are good independent reasons to accept the proposed revisions to our conceptions of lexico-syntactic structure and truth-conditional content.

The arguments I have set forth have indicated that it is at least questionable whether the approaches provide such motivation. In general, the conclusions reached in the preceding discussion all point to the idea that we should seriously question whether it makes sense to sacrifice our independently motivated conceptions of what constitutes lexico-syntactic structure and truth-conditional content in order to develop a compositional semantics for natural language along the lines proposed by the relevant approaches. That is, we are faced with a choice: either we accept the approaches offered at the cost of sacrificing certain intuitively and theoretically appealing notions about what constitutes linguistic structure and truth-conditional content, or we preserve the relevant notions at the cost of potentially having to sacrifice the idea that we can construct compositional semantic theories for natural languages.

Advocates of the positions discussed argue that the problem with going down the latter route is that once we abandon compositionality, we appear to abandon all hope of explaining the evidently systematic link between sentences and the interpretations they receive on occasions of use. That is, once we allow that our ability to interpret a sentence depends on something other than our ability to understand its compositionally determined linguistic content, it becomes difficult for the standard semantic theorist to imagine what else could possibly serve to forge the systematic link between sentential structure and truth-conditional content.

One might appeal to broadly pragmatic considerations here – i.e. considerations pertaining to particular kinds of background knowledge, interests, and practices that jointly motivate speakers to interpret sentences in the way they do on particular occasions of utterance. However, the standard semantic theorist will insist that the problem with this approach is that there is no obvious reason to think that such

considerations are answerable to the various constraints that appear to govern what interpretations speakers actually associate with sentences.

That is, they argue that whilst one can appeal to all sorts of reasons why one would *want* to interpret a sentence in a given way on a particular occasion of use, there are clearly limits on the ways speakers actually *do* interpret sentences that cannot be overruled by such pragmatic considerations. Ultimately, the standard semantic theorist is therefore led back to the view that the only plausible place to situate the relevant constraints on interpretation is in our understanding of the lexico-syntactic features of sentences themselves. Consequently, they cleave to the view that the truth-conditional contents of sentences are determined exclusively by their lexico-syntactic features, even if that means giving up certain independently plausible assumptions about what constitutes lexico-syntactic structure and/or truth-conditional content.

The findings in the thesis are clearly detailed in the relevant chapters. However, what they seem to demonstrate is that all of the accounts examined lack, in some way, the resources to explain how a standard semantic theory could fully support a truth-conditional externalist semantic theory. Each of the analyses appears to demonstrate that, in a number of cases, appeals to our linguistic competence alone may be insufficient to account for the interpretive phenomena. As we have seen, the standard semanticist is reluctant to appeal to resources outside of purely *linguistic* competence, principally because of fears of loss of systematicity. However, particularly in the context of discussions of the minimalist accounts, I have shown that that fear is unwarranted, at least in its extreme form, and an appeal to wider cognitive faculties need not involve a necessary loss of systematicity.

Of course, in one sense, this is a negative finding in that my analysis has shown gaps in the current prominent accounts, however, as just stated, the presence of such gaps and the reorientation of thinking about wider cognitive facilities expands the potential domain of inquiry. Indeed, it may be the case for those that wish to maintain and externalist, truth-conditional semantics, there is little choice but to look outside linguistic competence, as determined by the language faculty, for the explanation of the sorts of phenomena that I have examined.

The idea that we sometimes have to look beyond the capability of the language faculty in order to account for the relevant interpretive phenomena is not alien to all linguists. Pietroski is clear that we should not be surprised if fully accounting for truth takes us beyond the capabilities of linguistically determined meaning (Pietroski 2005b, p.254). None of this is to deny the value of the vast majority of standard semantic theorising. However, what I think we wish to avoid is the dogmatic effect that adherence to an exclusively *linguocentric* approach may cause. In particular, I think it is really important that researchers feel able to look beyond linguistic competence when the relevant phenomena warrant it, and do not feel, because of the inherent presuppositions of most standard semantic approaches, that they must account for all phenomena in the standard way.