



UNIVERSITAT<sub>DE</sub>  
BARCELONA

## Context and Compositionality: An Essay in Metasemantics

Adrian Briciu



Aquesta tesi doctoral està subjecta a la llicència **Reconeixement 3.0. Espanya de Creative Commons.**

Esta tesis doctoral está sujeta a la licencia **Reconocimiento 3.0. España de Creative Commons.**

This doctoral thesis is licensed under the **Creative Commons Attribution 3.0. Spain License.**

**University of Barcelona**  
**Faculty of Philosophy**  
**Department of Logic, History and Philosophy of Science**

**CONTEXT AND COMPOSITIONALITY**  
**AN ESSAY IN METASEMANTICS**

**ADRIAN BRICIU**

**Program: Cognitive Science and Language (CCiL)**

**Supervisor**  
**Max Kölbel**



## Contents

Introduction.....	9
1.The Subject Matter .....	9
2.The Main Claims.....	13
3.Looking ahead.....	14
CHAPTER 1: A General Framework.....	16
1. Semantic Theories: Aims, Data and Idealizations .....	16
2. Syntax .....	20
3. Semantics.....	25
3.1. Semantics and Syntax: Logical Form and Surface Form .....	25
3.2. Semantics: Total or Partial? .....	27
3.3 Assigning Meanings To Expressions: An Abstract Model .....	30
4.What is Meaning?.....	32
4.1 Context and Circumstances.....	32
4.2. Benefits of Double Indexing.....	38
4.3.Predictions of Truth for Utterances.....	42
4.4 Choice of Semantic Value .....	44
4.5 Extending the Framework: Choices in Modeling Content.....	46
5. Summing up.....	50
CHAPTER 2: Principles of Compositionality.....	55
1. Introduction .....	55
2. Formulating Compositionality .....	56
3. Compositionality and Semantic Values.....	59
3.1 Compositionality for Linguistic Meaning (Character Compositionality).....	60
3.2 Compositionality for Content .....	62
4. Summing up.....	75
CHAPTER 3: Context Sensitivity and Compositionality.....	77
1.Introduction.....	77
2.Semantic and Pragmatic Explanations.....	78
3.Data.....	81
4.Underdetermination .....	82
5.Underdetermination and Compositionality.....	87
6. How to Account for the Data: A Brief Geography of Positions.....	91
7. Truth Conditional Pragmatics .....	96
8. Pragmatic Intrusion and Weak Compositionality: The Model.....	101
9. The Wrong Type of Arguments from Context Sensitivity.....	103
10. Summing up.....	106
CHAPTER 4: Weak Compositionality and Systematicity.....	108
1. Introduction: Systematicity and Derivation of Truth-Conditions .....	108
2. Derivation of Truth-Conditions and Weak Compositionality.....	111
3. Open Texture and Systematicity.....	113
4. Weak Compositionality and Systematicity .....	117

5. Summing up.....	132
CHAPTER 5:.....	136
Weak Compositionality and The Productive Character of Natural Languages.....	136
1. Introduction .....	136
2. Compositionality and Productivity Facts.....	137
3. The Learnability Argument.....	138
4. The Explanatory Power of Compositionality.....	141
5. English and Inverted English.....	143
6. Weak Compositionality: Promises and Disappointments.....	147
7. Domestications.....	153
8. The Semantic Significance of Syntactic Rules.....	156
9. Summing up.....	162
CHAPTER 6: Conclusions .....	164
References.....	168

## Acknowledgments

There are many people and institutions to whom I have express thanks for providing help in the course of writing this dissertation. Institutionally, I have to thank the University of Barcelona and Agencia de Gestio D'Ajuts Universitaris I de Recerca (AGAUR) for funding my doctoral studies.

There are many people whom I have to thank for teaching me what I know, and/or for offering support, advice or encouragements. First I would like to thank the entire collective of the LOGOS Research Group at the University of Barcelona. I have learned a tremendous amount of stuff from them. I'd like to thank my supervisor Max Kölbel for long and fruitful discussions over the years on various topics in philosophy of language, as well as for many detailed comments and criticisms of earlier drafts of this dissertation, and, not least, for encouraging me to keep going. I'm also indebted to my colleagues David Rey, Dan Zeman, Fiora Salis, Andrei Moldovan, Oscar Cabaco, Sana Hirvonen, John Horden, Bryan Pickel, Stephan Torre, István Aranyosi, Lucian Zagan, Romina Zuppone and Lucía Prieto Castrillo.

Earlier drafts of some of the chapters were written while I was on two visiting Institut Jean Nicod in Paris and Stockholm University. I would like to thank Peter Pagin, François Recanati, Dag Westerståhl, Sara Packalén, Marie Guillot, Emma Wallin and Michael Murez for discussions on various subjects covered in this thesis, as well as for many useful suggestions.

I would like to thank my partner, Lucía, not only for helpful discussions but also for all the support, encouragement, advice and comfort she provided, in the absence of which this dissertation wouldn't have taken shape.

## Summary

My general aim in this dissertation is to analyze two intersecting topics. One is what constraints the principle(s) of compositionality place on semantic theories, and the second is whether context-sensitivity endangers the project of formal semantics. More precisely I will look on what constraints different principles of compositionality place on the treatment of context-sensitivity and I will focus primarily on whether the claim that natural language sentences exhibit pervasive and radical context sensitivity (that allegedly cannot be treated can be handled by fixing the values for a fixed set of contextual parameters) can be accommodated within a compositional semantics.

I approach this debate from a slightly different angle than most authors. Most authors are concerned with whether the arguments that supposedly establish the existence of such context-sensitivity are correct or not. I, on the other hand, approach it from a different angle: assuming (or conceding) that there is pervasive radical context-sensitivity, is this incompatible with the claim that natural language semantics are compositional? The most common answer is that indeed there is incompatibility between this kind of context-sensitivity and some forms of compositionality. To give a precise answer to this question I'll distinguish three principles of compositionality, where each of them places different constraints on semantic theories. I'll show that radical context sensitivity this is incompatible with a strong version of compositionality but not with a weaker version of compositionality. The strong version claims that the semantic value of a complex expression relative to a context is a function of the semantic values of its constituents (at that context) and of its syntactic structure. The weaker version claims that the semantic value of a complex expression at a context is a function of the semantic values of its constituents (at that context) its syntactic structure and of the context itself. Ultimately, the question I address can be formulated in the following way: is the contribution of the context to the determination of truth-conditions of sentences at contexts restricted to that of fixing the values of context-sensitive expressions (or other syntactic material) present in

the logical form? The stronger version of compositionality forces us to give a positive answer, while the weaker version of compositionality permits a negative answer.

In the final chapters of the dissertation I offer several arguments against the weak version of compositionality. The main contention is that weak-compositionality fails to deliver the explanatory benefits that we expect compositionality to deliver. If the arguments are correct this limits the possible moves a theorist can take in handling context-sensitivity.





# Introduction

In this introduction I'll give a general and non-technical overview of the dissertation. First, I'll give a brief and non-technical presentation of the subject matter that this thesis deals with, and then I'll briefly present my central claims. I'll end the introduction with a brief overview of the organization of the thesis.

## 1. The Subject Matter

It is fair to say that, until recently, philosophers of language could be divided in two camps: those who believed that the fundamental meaning-properties of natural languages and that linguistic competence with them can be modeled with the tools of formal logic (call them formal semanticists) and those who were skeptics about the prospects of such project (call them simply the skeptics). Much of the subject matter of this thesis is the result of confrontation between these two camps.

The guiding light of formal semanticists is the belief that the meaning-properties of natural languages can be described (or modeled) with the help of formal languages; that there is no essential difference between formal languages and natural languages in the sense that for any natural language we can build a theory by means of which we can deductively establish the meaning of any sentence of that language. There are many ways to implement this program but most formal semantic theories share several basic ideas.

One of them is the idea that all there is to the meaning of a sentence (at least as far as modeling is concerned) is its truth-conditions. This originates with a basic insight about how we should handle meaning. The insight is that to know the meaning of a sentence is to know the conditions under which the sentence is true and the conditions under which it is false. English speakers (just like speakers of any other natural language) have this manifest ability: when provided with an English sentence and with a range of scenarios, ideally, they

can divide these scenarios in two classes: those of which the sentence is true and those of which the sentence is false. If we take these kinds of abilities to reflect crucial aspects of meaning and of linguistic competence, then we're provided a solid footing for theorizing about meaning in natural languages. But being so general this idea leaves a lot of maneuvering space to theorists on how exactly to model meaning. For example, one dimension of this maneuvering space is that theorists can decide between a semantics that pairs sentences with their truth-conditions or a semantics that pairs sentences with rules or conventions that determine their truth-conditions for every occasion of use.

A second basic idea shared by virtually all formal semantic theories is the principle of compositionality which states that the meaning of complex expressions is determined by the meaning of their constituents and their syntactic structure. Now this formulation is quite general and somewhat imprecise in the sense that it can have several different (i.e. non-equivalent) specifications. In fact, it is fair to say that the principle of compositionality is quite often approached more like a tacit assumption: more often than not semanticists do not bother to give it a precise formulation and often they simply assume that their theories conform to it (or can be made to conform to it). Given that the principle can be made precise in different ways then it is no surprise that different semantic theories will satisfy different versions of compositionality (or rather different principles of compositionality). Of course, any theorist is entitled to their own principle of compositionality but any principle of compositionality worth its name must deliver certain explanatory benefits. And most theorists consider compositionality to be an essential part of (a) any plausible account of the productive features of natural languages (i.e. of how speakers with limited cognitive means can understand any of an indefinite number of meaningful sentences) and of (b) any account of the systematic features of natural languages (i.e. of how meaning of sentences is systematically determined from the basic meaningful units of the language). Then, rather obviously, we will judge any principle of compositionality depending on whether it can deliver on its promises. The different principles of compositionality can vary along many dimensions. For example, it is often emphasized that the principle (in its general formulation) can be made precise only together with an explicit theory of meaning and of syntax, that is, with a precise specification of what meanings are and of how expressions combine. Then, as I'll explain in detail in

chapter two, if a theorist prefers a semantics that assigns truth-conditions (as meaning) she will work with a different principle of compositionality than if she prefers a semantics that assigns as meaning conventions (or rules) that determine truth-conditions. The reason is context-sensitivity and the role that context plays in the determination of meaning. There is some sense in which some expressions of natural languages express different meanings on different occasions of use. English expressions (and their correspondents in other languages) like "I", "here", "there", "now", "yesterday", "that", "to the right" are the most obvious and uncontroversial examples. It is quite obvious that what makes true or false a sentence that contains context-sensitive expressions can vary from occasion of use to occasion of use. Context-sensitivity was one of the first places where it looked like formal semantic theories face serious difficulties: for how can a theory establish in a systematic manner the truth-conditions of any sentence (of a given natural language) if those truth-conditions vary with the occasion of use? The solution was to take the linguistic convention associated with simple expressions to be or to provide a rule which, given an occasion of use, determines the content of that expression on that occasion. For example, the idea is that we can take the linguistic convention associated with "I" to be, or to provide, the rule that on any occasion of use "I" refers to the speaker. If context sensitivity can be treated along these lines then there is hope that it can be accommodated by formal semantic theories. If the theorist opts for a semantics that assigns truth-conditions, then truth-conditions will be assigned to sentences at contexts of use, and the truth-conditions of a sentence (relative to a context) is a function of the content of its constituents at that context and of its syntactic structure. This is what I'll call strong compositionality for content. If she opts for a semantics that assigns linguistic meaning that on any given context of use determines the truth-conditions a sentence has at that context, then she will say that the linguistic meaning of sentences is a function of the linguistic (i.e. conventional) meaning of their simple constituents and of its syntactic structure. I'll call this the principle of compositionality for linguistic meaning.

Skeptics, on the other hand, doubt that all context-sensitivity can be treated in this way and, therefore, that the project of formal semantics is worth pursuing. According to them context-sensitivity is much more radical and pervasive in natural languages than what formal semantics can handle in the sense that much of it cannot be handled in an analogous

way with indexicals: that is by fixing the value of a limited set of contextual parameters (such as the speaker, time and place of utterance and sequence of salient objects). Their argumentative strategy is the following. They put forward a series of simple but imaginative examples that consist, usually, of an English sentence and a couple of scenarios. They elicit the intuition that in order for that sentence to be true of both scenarios, different states of affairs must obtain in those scenarios. Furthermore they argue that the difference in the truth-conditions of the sentence cannot possibly be the result of fixing the values of a limited set of contextual parameters. The difference in truth-conditions, so they argue, can be explained only by appeal to the intentions, interests and practical purposes that the conversational participants have at those respective scenarios. Intentions, interests and practical purposes are, in some sense, formally intractable and out of the reach of semantic theories. Then any theory that takes into account just the linguistic properties of expressions will fail to derive truth-conditions for sentences, for truth-conditions are severely underdetermined by the linguistic meaning of sentences. According to skeptics, even for the simplest sentences various pragmatic considerations about the conversational partners interests, concerns and intentions will select highly specific (and sometimes perspectival) aspects of the world that will make a sentence true or false. If skeptics are right, then strong compositionality is not true of English (and other natural languages): for there are sentences that have different truth-conditions on different occasions of use, although none of their constituents is context-sensitive (and hence none of their constituents vary their contents across occasions of use).

Lately though, many other authors have argued that there is no incompatibility between the acceptance of such pervasive and radical context-sensitivity and the possibility of systematic semantics, where by systematic semantics is understood a theory that systematically derives truth-condition for sentences at contexts. Supposedly weak compositionality will help the theorist handle pervasive and radical context-sensitivity in a systematic manner. According the weak compositionality for content, the truth-conditions of a sentence (relative to a context) are a function of the content of its constituents (at that context) of its syntactic structure and of the context itself.

## 2. The Main Claims

My main claims concern the notion of weak-compositionality. Succinctly put, I'll argue that it fails to deliver its expected explanatory benefits. It is argued that weak compositionality will help theorists handle pervasive and radical context-sensitivity in a systematic manner. I'll give several arguments that will show that that is not the case: if a theory accepts pragmatic intrusion into the derivation of truth-conditions even if the theory is weakly compositional it *can* still fail to determine a systematic procedure for derivation of truth-conditions. Moreover, as I'll discuss in some detail, there are no reasons to believe that there is a systematic procedure that will determine for every sentence-context pair the pragmatic factors that are relevant for the determination of truth-conditions of that sentence at that context. In fact there are plenty of good reasons to believe just the opposite. I will also argue that weak-compositionality fails to explain the productive features of natural languages. I'll give several arguments in this direction. First, I'll argue that in a weakly-compositional model speakers have to learn one by one all of the indefinitely many ways in which context can affect the truth-conditions of a given sentence. Secondly, I'll argue that no weakly-compositional theory that allows for pervasive and radical context-sensitivity can be finitely stated. This argument concerns the form of semantic rules that such a theory must employ. More precisely, I'll argue that the semantic rules of a theory that pairs sentences in contexts with their truth-conditions, and that allows context to freely contribute to truth-conditions, cannot be finitely stated. Moreover, I'll also argue that no further constraints can be placed on such theories that can alleviate these problems. Thirdly, I'll raise some worries about the principle of weak-compositionality, independently of whether it is satisfied by semantics that allow context to freely contribute to truth-conditions. That is, I'll consider the possibility of semantics that satisfy weak-compositionality but that rejects that context contributes freely to truth-conditions. Such semantics are not open to worries about failure to explain the productive features of natural languages, but, as I'll argue, they presuppose an implausible account of the semantic significance of syntactic rules and of what knowledge of syntactic rules brings to the economy of knowing the meaning of complex expressions.

### 3. Looking ahead

In the first chapter I'll present a general semantic framework that states clearly the most important assumptions and tenets of (most) formal semantic theorize. It is a framework that permits to give a precise formulation of the problems surrounding context-sensitivity, as well as it permits to give a precise formulation of different principles of compositionality and to state clearly the relations among them. In the second chapter I'll do exactly that: I'll give a precise formulation of different principles of compositionality and of the relations among them. In the third chapter I'll discuss the skeptic challenge. I'll start with a brief overview of the data and the arguments put forward by the skeptics, and then I'll give my own proposal about how we should understand the skeptic challenge and how it relates to different versions of compositionality. I'll end the chapter with a brief presentation of truth-conditional pragmatics and of how weak-compositionality makes room for pragmatic intrusion in the determination of truth-conditions. In the fourth chapter I'll analyze on whether weak compositionality can deliver some of its promises. I'll show that weak-compositionality will fail to reconcile the kind of pervasive and radical context-sensitivity that skeptics take to threaten the project of formal semantics with systematic derivation of truth-conditions. More precisely, I'll show that theories that accept pragmatic intrusion into the determination of truth-conditions although they can satisfy weak-compositionality, they fail to specify a systematic manner to derive truth-conditions. In the fifth chapter I'll argue that weak compositionality fails to explain the productive features of natural languages. I'll consider one possible rejoinder and I'll argue that even accepting it, weak-compositionality still presupposes an non-standard and implausible account of the semantic significance of syntactic rules. I'll end the thesis with a discussion of the constraints that the other two versions of compositional (introduced in chapter two) place on semantic theories and on accounts of context-sensitivity.





# CHAPTER 1: A General Framework

## 1. Semantic Theories: Aims, Data and Idealizations

In this chapter I'll present a general semantic framework. This will help pin down some important things: what semantic theories are about, what assumptions and idealizations they make, what is the proper shape of semantic theories, as well as it will help identify a precise target for skepticism towards formal semantics.

What do semantic theories seek to explain and what sort of questions do they try to answer? A brief answer, which suffices for now, is that semantic theories are in the business of giving an abstract explanation of linguistic competence with natural languages: of what is it to speak and understand a natural language. It achieves this by giving formal characterizations of natural languages that can serve as explanatory models of linguistic competence. These formal models will represent certain linguistic properties and relations that individuate a natural language: grammaticality, ambiguity, synonymy, analyticity (i.e. truth in virtue of meaning), contradiction, and so on. We see this better if we look at what sort of data semantic theories try to predict, and the evidence that they must be answerable to.

There are two types of data that semantic theories must explain and must be answerable to. One are *intuitions (or judgments) about the truth-values* of natural language sentences relative to actual and possible scenarios. Speakers know whether a sentence is true or not if used in such and such occasion and the world is such and such - where the occasion of use can be either actual or possible, and the world can be either actual or possible. A second type of data that semantic theories must explain is the *productive character of natural languages*: speakers of natural languages can produce and understand complex meaningful expressions that they have never encountered before, and they can understand an open-ended number of such expressions, more than they could ever learn in a one-by-one manner.

Let me explain the first claim a bit. Competent speakers of a language know which

expression are grammatical and which are not, which sentences are ambiguous (and what the ambiguity consists in), which expressions mean the same and which not, they know that some sentences entail others, while some sentences are inconsistent with others, they have the intuition that some sentences are true relative to some actual or possible scenarios, while others are false relative to the very same scenarios, and so on. These linguistic properties (synonymy, ambiguity, etc) and speakers' knowledge of them are facts that semantic theories need to explain. There must be something in speakers' linguistic behavior that is fairly unproblematic and accessible to theorists and that reliably reflects this knowledge. One fruitful idea is that intuitions about correct usage of expressions might just provide what is needed: competent speakers have the ability to tell what would make an utterance correct or incorrect, and they are easily accessible to theorists. Then intuitions about appropriateness or correctness of utterances are data that semantic theories aim to predict. Now, there are many dimensions along which utterances of natural language sentences can be judged to be appropriate or correct. For reasons that will be presented in the following sections one dimension of correctness, at least when it comes to declarative sentences, is of particular interest to the semanticists. Utterances of sentences are judged to be correct (or appropriate) if they are true relative to actual or possible states of affairs. Then, intuitions about the truth-values of natural language sentences relative to actual or possible scenarios are data that semantic theories aim to predict. These intuitions are fairly accessible to theorists and they are taken to reliably reflect speakers' competence with the sentences of the language. Moreover, it is enough that a theory gives truth-value predictions for sentences of a language in order to predict analyticity, synonymy, entailment, in that language. (That is to say that from truth-value predictions the theorist can derive predictions about analyticity, synonymy, entailment.)

The semanticist's data, then, comes from language use. But language use quite often is messy, vagarious and loose. The daily usage of language, in its brute state, might be of little use to the semanticist: speakers don't always use full sentences, they change their plan in mid-phrase, sometimes they use words in idiosyncratic ways (e.g. malapropisms), they often pepper their speech with metaphors and other creative uses of words, they help themselves of non-linguistic props (e.g. "uhm", "huh?") or they use sentences in apparently incongruous ways (e.g. like when one utters "John is a fine friend" on one occasion to

communicate that John is a fine friend, and utters the same sentence on a different occasion to communicate the opposite). Before any proper semantic work can be done, the job of the theorist is to filter out the data that semantic theories need to explain from the noise of language use. And here several idealizations come in.

First the theorist will distinguish *linguistic competence* (what speakers know when they know a language) from *performance* (the actual use of language in concrete situations). This will help her separate those aspects of language use that are the result of linguistic competence from those that are the result of performance mechanisms. Theorists will focus on somewhat idealized speakers, in a completely homogenized linguistic community, that are not affected by irrelevant conditions like memory limitations, distractions, limitations of attention span, articulatory impairments and so on. Only under this idealization, uses of language are a direct reflection of linguistic competence - Chomsky (1965) is the *locus classicus* for this distinction.

I said above that the data that semantic theories must predict are intuitions about truth-values of sentences relative to actual and possible scenarios. But this is too general and somewhat imprecise. Speakers can judge the literal truth of an uttered sentence, the truth of the message conveyed by that utterance, or the metaphorical truth of an utterance. For example, an utterance of “I watched the election debate a million times” might be literally false, but the message conveyed by it on a given occasion might be true. What semantic theories are interested in predicting are intuitions about the truth-values of literal uses of sentences (i.e. intuitions about literal truth) – I'll explain why this is so in the subsequent sections. Again, there is some filtering work to be done before any proper semantic work can start. So the theorist must offer a way of distinguishing those intuitions (or judgments) about truth-values that are about the literal uses of sentences from those that are about their non-literal aspects (e.g. that are about the truth of what was conveyed by an utterance of that sentence) – Grice (1989) is the *locus classicus* for this distinction. This is not to suggest that drawing this distinction is easy or that it is just a methodological point. In fact, where to draw the distinction is one of the most controversial topics in philosophy of language and linguistics. As I'll discuss in chapter 3 there is still a lot of controversy concerning intuitions about truth-values of many sentences: are they intuitions about their

literal truth or intuitions about the truth of what is conveyed by utterances of those sentences? The idea is that, in principle, we can decide in a non-arbitrary fashion for any intuition about truth-values whether it is or not about the literal truth of the sentence.

It is clear from what I said that the data that semantic theories need to explain are somewhat idealized and the result of theoretical filtering. Semantic theories make predictions about these idealized data, abstracted from the noise of daily use, and any skepticism towards the predictive power of formal semantic theories should take this into consideration.

In a nutshell: formal semantic theories are in the business of modeling natural languages, and linguistic competence with them, with the tools of formal logic. Basically, semantic theories offer formal characterizations of natural languages that can serve as explanatory models of linguistic competence with those languages. The model-language constructed by the theorist is an abstract object and its semantic properties are the result of theorist's stipulations. Then the theorist must then decide whether the model-language as described by her theory corresponds to, or is used by any group or population of language users. This is where the model makes contact with linguistic reality and where formal semantic theories get their empirical content. (As I'll explain in subsequent sections, the predictions of the theory are about sentences of the model language, in the first place, and are about natural language sentences only given some bridge-laws that connect the two). Given that semantic theories are supposed to capture both intuitions about truth-values and the productive character of natural languages they will be strongly constrained by the need to explain the possibility of productivity. Then the model-language will include at least (a) a finite vocabulary pairing atomic expressions with semantic values, (b) a finite set of syntactic rules which build complex expressions out of the simpler ones, (c) a finite set of rules that determine the semantic value of complex expressions as a function of the semantic value of constituents and their syntactic structure. Finally, it should be stressed that formal semantic theories give an abstract explanation of linguistic competence, of what one knows when she knows the meaning of sentences and how is it possible that linguistic competence is productive. They do not speculate about the causal mechanisms that underlie or implement linguistic competence.

The plan of this first chapter is the following. First I give a short description of how we can think about syntax. After that I'll give a short description of how we can think of meaning (semantic value), and I present this within a widely used semantic framework inherited from Kaplan (1989) and Lewis (1980). At the end of the chapter I briefly sketch several types of semantics that can be given inside this framework. This will help with several things. First it will help in giving a precise definition of the principle of compositionality, and it gives a precise formulation of problems that stem from context-sensitivity.

The most common formulation of the principle of compositionality found in the literature (what I will call the traditional formulation) is quite general and not free of ambiguities:

*the meaning of a complex expression is a function of the meaning of its constituents and of the way the constituents are combined.*

To put more flesh on this principle I'll have to say what meanings are and what ways of combining constituents are. This is what I'll do in the following sections, and I'll start with ways of combining expressions.

## **2. Syntax**

In characterizing a language theorists usually distinguish syntax and semantics. Roughly, the syntax of a language states what the expressions of the language are, while the semantics says what the expressions of the language mean. In the next sections, I will introduce a standard way of theorizing about the syntax and semantics of languages.

A general discussion about what syntax is, will help to specify several important notions used in the formulation the principle of compositionality: those of *simple (atomic) expressions*, *immediate constituent* and *syntactic structure*.

What is syntax? The syntax of a language states what the expressions of that language are. This is done by specifying a set of expressions and specifying a set of rules

that determine which combinations of expressions are also expressions of the language. In other words, the syntax of a language is a structure that consists of a set of expressions and a set of rules of combination for those expressions (grammatical rules for natural languages; formation rules for artificial languages). Syntactic rules state how expressions of the language can combine to form more complex expressions. We can think of syntactic rules as functions defined on the set of the expressions of the language and with values in same set.

This is very versatile and general way of stating what syntax is. It is general enough to be helpful in discussing the syntax of both formal languages and natural languages. We can follow these general lines and state the syntax of a formal language that can serve as a model for natural languages. To do this, we need a set of expressions and a set of syntactic rules combining those expressions. Take  $E$  to be the set of all expressions of a language  $L$ , take  $A$  to be the set of all simple (atomic) expressions of  $L$  where  $A$  is a subset of  $E$ , and take  $F$  to be the set of its syntactic rules. Syntactic rules combine simpler expressions into more complex ones. Thus, as a first approximation one natural way to define syntactic rules is as functions from members of  $E$  to members of  $E$ : every syntactic rule  $\gamma$  that belongs to  $F$  is a diadic function from elements of  $E$  to elements of  $E$ <sup>1</sup>. Every member of  $E$  is either a simple expression (i.e. member of the subset  $A$ ) or is the value of an operation of  $F$ . No other expression is a member of  $E$ . In other words  $E$  is closed under operations in  $F$ . We say that an expression  $e$  is simple (atomic) if it belongs to  $E$  and there is no function  $\gamma$  that belongs to  $F$  such that  $e$  is the value of  $\gamma$  for any two arguments. Otherwise, if  $e$  belongs to  $E$  and there is a  $\gamma$  of  $F$  such that  $e$  is the value of  $\gamma$  for some arguments then  $e$  is a complex expression. We also say that an expression  $e$  is an *immediate constituent* of a more complex expression  $e$ -complex if both  $e$  and  $e$ -complex belong to  $E$  and  $e$ -complex =  $\gamma(e, e^*)$  for some function  $\gamma$  that belongs to  $F$  and some additional argument  $e^*$  that belongs to  $E$ .

It is a common observation that some expressions of natural languages are of different kinds (that we call *grammatical categories*) while others are of the same kind. For

---

<sup>1</sup>I take syntactic functions to be diadic functions, while most commonly in the literature the number of arguments remains unspecified - they are taken to be as *n-adic* (where  $n > 1$ ). I do it mainly for reasons of presentation: it simplifies the exposition later on, and nothing gets lost. On the other hand, allowing only diadic composition is a significant limitation. But everything that I say regarding syntax, semantics and compositionality can be reformulated, with the same results, taking syntactic operators to be *n-adic* functions (where  $n > 1$ ).

example, the English expressions “man” and “woman” are of the same category but of different category from “walk” and “talk” which in their turn are of the same category, and all four are different from “if” and “when”, which in their turn are of the same category. Thus, when we give the syntax of a natural language we can categorize every expression of  $E$  into grammatical categories. Every expression of  $E$  is a member of a particular grammatical category: noun, verb, determiner, adverb, adjective, preposition, verb phrase, noun phrase, determiner phrase, sentence and so on. Let  $Cat$  be the set of all grammatical categories of that language, then for each category  $c_x$  such that  $c_x$  belongs to  $Cat$ ,  $E_{c_x}$  is the set of all expressions of that category. For any expression  $e$  such that  $e \in E$  there is a grammatical category  $c_x$  such that  $e \in E_{c_x}$ . Then syntactic rules don't have to yield a value for any pair of expressions of  $E$ , but we can take syntactic rules as functions defined over expressions of certain categories and with values in expressions of certain categories. In other words, each syntactic rule can be specified in terms of the categories of its arguments as well as the category of its value.

For the purpose of illustration I will describe the syntax of an artificial language  $L^*$  that consists of three simple expressions: “John”, “Mary”, “runs”, one syntactic rule, and expressions generated by that rule. One can think of this language as identical with a fragment of English, and thus see how the above framework can be used to describe the syntax of English.

The syntax of  $L$  is a structure  $\langle E, F \rangle$  where  $E$  is the set of expressions of  $L^*$  and  $F$  the set of syntactic operation defined on  $E$  and with values in  $E$  such that:

- a. The set  $A$  of simple expressions of  $L^*$ ,  $A = \{\text{John, Mary, runs}\}$ , where  $A \subseteq E$
- b. The set  $Cat$  of grammatical categories of  $L^*$ ,  $Cat = \{\text{NOUN, VERB, SENTENCE}\}$
- c. The set of syntactic operations  $O = \{\text{concatenation}\}$
- d. A function that assigns to each expression of  $E$  a category from  $Cat$  and the sets of category-indexed expressions generated by that function:  $E_{\text{VERB}} = \{\text{runs}\}$ ,  $E_{\text{NOUN}} = \{\text{John, Mary}\}$ ,  $E_{\text{SENTENCE}} = \{\text{John runs, Mary runs}\}$
- e. The set  $F$  of syntactic rules of  $L^*$ :  $F = \{\gamma\}$  where and  $\gamma$ : every concatenation of an expression of category  $\text{NOUN}$  with an expression of category  $\text{VERB}$ , in this order,

results in an expression of category SENTENCE.

This model tells that some combinations of expressions of  $L^*$  are themselves expressions of  $L^*$  (for example “John runs” and “Mary runs”) while other combinations are not expressions of  $L^*$  (for example “runs Mary”, “runs John” or “John Mary”). This can work as a model for a very small fragment of English. Not every combination of English expressions is itself an English expression, “runs Mary” is not an English expression (we say it is ungrammatical) while “Mary runs” is (we say is grammatical).

Instead of listing all grammatical categories we can do the same with a small adjustment, and consider that there are only a very small number of *basic grammatical categories*, for example, the category of name (N) and another is the category of sentence (S). All other grammatical categories are *derived*. For example, there is the category of expressions that combine with names N to give expressions of category S (we write this category S/N). These are all grammatical categories that  $L^*$  has. Now, when we specify A (the set of simple expressions of  $L^*$ ) we assign to each expression a category. The categories of expressions members of A are either basic or derived. No member of A is of category S, some are of category N, and, in order to be able to generate complex expressions of category S, some must be of the derived category S/N. If “John” and “Mary” are of category N then, “runs” must be of category S/N. Then  $\gamma$  can be specified as: any concatenation of an expression of category N with an expression of category S/N, in that order, results in an expression of category S. (Giving this rule explicitly is a bit superfluous, because the rule is intrinsic in how we define the derived category S/N - technically we could do without specifying explicitly syntactic rules like  $\gamma$  and just lay down a general rule for generating grammatical categories).

This again is very versatile and it is easy to see how it can be used for languages more complex than  $L^*$  (or for larger fragments of English). We can use the set of basic grammatical categories and the procedure employed above to derive far more grammatical categories than just S/N. For example we can derive the category of expressions that combines with names N to give the category of S/N (we write it as (S/N)/N)<sup>2</sup>.  $L^*$  doesn't

---

<sup>2</sup> S/N corresponds to the traditional grammatical category of Verb-Phrase, while (S/N)/N to the traditional grammatical category of Intransitive Verb. Whether we use traditional grammatical categories or categories like S/N and (S/N)/N, when we specify the elements of A is irrelevant. For this point see Lewis 1972: 193 and E. Bach



contain this category. But if we want to extend  $L^*$  (or consider a larger fragment of English) with the expression “loves” this category is needed. And, again, we can extract the corresponding syntactic rule: if an expression of category  $N$  is prefixed by concatenation with an expression of category  $(S/N)/N$  the result is an expression of category  $S/N$ . Obviously we can easily extend this further on. In fact, in this way, we can specify for every expression of a natural language its category. To make it clear, we can write down the syntax of  $L^*$  in the following way:

- a. The set  $A$  of simple expressions together with a function that assigns to each member of  $A$  a category from set  $C: \{John_N, Mary_N, runs_{S/N}\}$ ,
- b. the set  $Cat$  of all grammatical categories  $\{S, N, S/N\}$
- c. the set  $O$  of syntactic operations  $O = \{\text{concatenation}\}$
- d. the set  $F$  of syntactic rule,  $F = \{\gamma\}$  where  $\gamma$ : every concatenation of an expression of category  $N$  with an expression of category  $S/N$ , in this order, results in an expression of category  $S$ .

This is still very general, but it gives an idea and a blueprint of how the syntax of natural languages can be written down. Once we move from  $L^*$  (or very small fragments of English) to natural languages like English (or, at least, to large fragments of English) unexpected complications might show up, but it is generally accepted that they can be dealt with in this framework<sup>3</sup>.

This general specification of syntax gives us the tools to talk about expressions and syntactic rules. This matches with our pre-theoretical insights that languages are made up of expressions and rules that combine those expressions. But more importantly the ability to talk about expressions and syntactic rules is important when it comes to giving a precise definition of the principle of compositionality.

---

1988: 23

<sup>3</sup> E.Bach (1988) gives a general overview of different theories done along these lines.

### 3. Semantics

If a language consists of expressions and their meanings, and syntax specifies the expressions of the language, semantics assigns meanings to those expressions. So, we can think, in very general terms about the semantics of a language as a function from the expressions of the language to the meanings of those expressions. In other words, semantics is a function from members of E into a set M of meanings. I will say that semantics is a function **I** that interprets expressions in E.

It is very important what entities meanings are (i.e. what entities semantics assigns as meanings) but for the moment I will not specify that here. In the second part of the chapter I will consider several kinds of entities that can play the role of meaning.

#### 3.1. Semantics and Syntax: Logical Form and Surface Form

Natural languages, like English, contain syntactically ambiguous expressions. For example the English sentence

(3) “Mary saw the boy with binoculars”

has two distinct meanings. For one, it can mean that Mary saw the boy who had binoculars. It can also mean that Mary saw the boy with the help of binoculars. Similarly, the sentence

(4) “Every boy loves some girl”

has two distinct meanings. For one, it can mean that there is a girl such that every boy loves her. It can also mean that for every boy there is a girl that he loves, and it is possible that each boy loves a different girl.

A closer look at the first example, reveals that in the first reading of (3) the noun phrase “the boy with binoculars” is the direct object of the verb “saw”. While in the second reading of (3), there is no noun phrase “the boy with binoculars”, but the phrase “with

binoculars” is an adverbial phrase modifying the verb phrase “saw the boy”. If so, it means that the two readings of (3) have different syntactic structures.

For example, we can represent the first reading as:

(3a) [Mary<sub>N</sub> [saw<sub>V</sub> [the<sub>DET</sub> [boy<sub>CN</sub> [with binoculars<sub>AdjP</sub>]<sub>NP</sub>]<sub>NP</sub>]<sub>VP</sub>]<sub>S</sub>

and the second reading as:

(3b) [Mary<sub>N</sub> [[[saw<sub>V</sub> [the<sub>DET</sub> [boy<sub>CN</sub>]<sub>NP</sub>]<sub>VP</sub> [with binoculars]<sub>AdvP</sub>]<sub>S</sub>

Similarly, for the second example; (4) (I give it in less detail)

(4a) [Every boy]<sub>NP</sub><sup>i</sup>[[some girl]<sub>NP</sub><sup>j</sup>[\_]<sub>VP</sub><sup>i</sup> [loves \_]<sub>VP</sub><sup>j</sup>]]<sub>S</sub>

(4b) [Some girl]<sub>NP</sub><sup>j</sup> [every boy]<sub>NP</sub><sup>i</sup>[\_]<sub>VP</sub><sup>i</sup> [loves \_]<sub>VP</sub><sup>j</sup>]]<sub>S</sub>

The two readings of (3) require the distinct syntactic structures given in (3a) and (3b). It must be then that the semantically relevant syntactic structure of (3) is not its immediately apparent structure but either of the two distinct structures in (3a) or (3b).

Thus, it is helpful to adopt the following hypothesis: expressions in natural languages have what we might call a *surface form* (i.e. their immediately apparent syntactic structure) and a *logical form* (i.e. a syntactic structure that is relevant for their semantic interpretation). In order to account for syntactic ambiguity we need to posit and distinguish a level of syntactic structure that is different from the immediately apparent structure of sentences<sup>4</sup>. The logical form of an expression is the *disambiguation* of its surface form and is the syntactic structure that gets interpreted semantically. The surface form of an expression determines how that expression gets uttered (that is how it is pronounced or written). The surface form of an expression is the result of *transformations* of the logical form of the expression. For example, the displacement of the quantificational expression “some girl” from the position in (4b) to its audible position in the surface form of (4), is such a transformation.

We will say that (3) is not a member of E but the structures in represented in (3a) and (3b) are members of E. Again, for our purposes we can say that the members of E are

---

<sup>4</sup> There are plenty of data, besides syntactic ambiguity, that such a hypothesis helps account for. See Pietroski 2006 and Chomsky 2000 for an overview of the data.

not expressions of natural languages (say expressions of English) but *disambiguated expressions* (or expressions of a *model-language*). Thus, E contains two distinct expressions (3a) and (3b). Both expressions get semantically interpreted, (3a) as saying that Mary saw the boy who had binoculars and (3b) as saying that Mary saw the boy with the help of binoculars. Since both (3a) and (3b) undergo transformations into (3), the two expressions share the same surface form.

Since the logical form an expression is the syntactic structure relevant for semantic interpretation, in the rest of the thesis when I will be speaking about syntactic form I will be speaking about the logical form of expressions and not about their surface form. Similar considerations apply to lexical ambiguity. We will say that E contains two different expressions  $bank_1$  and  $bank_2$  corresponding to the two different meanings that the ambiguous surface expression “bank” has in English that share that share the same surface form but that get different semantic interpretations.

Given that we want semantic theories to make predictions about natural language sentences and the data available are intuitions about sentences like (3) and not about anything like (3a) or (3b) we will have to assume that expressions in the logical form like (3a) and (3b) are representations of surface form expressions like (3).

### **3.2. Semantics: Total or Partial?**

Another issue that needs clarification is whether semantics is a total or a partial function over expressions. Is it the case that every expression in E is a meaningful expression?

Some semanticists (e.g. Montague 1974) prefer to have a semantics that does not assign a meaning to every simple expression in E. For example some prefer to have some expressions, like “and”, “not”, “possibly” introduced directly by rules. These are mainly truth-functional operators to which no syntactic category is assigned in E. But “not”, “and”,

“possibly” are, nevertheless, expressions that belong to English, and most other natural languages have corresponding expressions. For this reason, other semanticists (e.g. Heim and Kratzer 1998) prefer a semantics that assigns meaning to every simple expression in E, including to those that function as truth-conditional operators. For the purposes of this thesis I will also assume that a semantics assigns meaning to every simple expression of E.

A different question is whether a semantics should assign meaning to every complex expression of E. The most common answer is 'no': natural languages like English contain meaningless expressions. Some authors (Chomsky 1957, Carnap 1928)<sup>5</sup> take

(5) “Colorless ideas sleep furiously”,

to be such an expression. If so **I** should be taken to be a partial function on E. There are different motivations for such a claim depending on what one’s favorite theory of meaning is. But probably the most common motivation runs along the following lines: to give the meaning of a sentence is to specify its truth conditions (i.e. what has to be the case such that that sentence is true). This is not possible for sentences like (5) since we cannot specify, nor conceive what has to be the case such that (5) is true, therefore (5) must be meaningless.

Other authors (Davis 2003, Magidor 2009) take (5) to be both a member of E and meaningful although they take its meaning to be anomalous or absurd, since abstract things like ideas are not the things that have colors, nor are they things that sleep (see Davis 2003: 260-261). If (5) is a member of E and meaningful, and the same holds for all expressions of the language (i.e. there are no expressions that are members of E but are meaningless) then **I** is not a partial but a total function. Semantically we should treat (5) on a par with other expressions (like “I am hungry”, “The moon is a satellite” and so on) and their anomaly should be explained in pragmatic terms<sup>6</sup>.

---

<sup>5</sup>The example is due to Chomsky (1957). Carnap (1928) gives the following examples “This rock is sad”, “This triangle is virtuous” and he offers what I just called the standard justification for taking such expressions to be grammatical but meaningless. “[the constituents of the sentences] are conjoined as their grammatical characters require but not as their meanings do. ...[i]f a statement does not express a (conceivable) state of affairs then it has no meaning; it has only apparent meaning” 1928/2003: 325)

<sup>6</sup>There are different motivations for such a position. To mention the most common: (5) can be embedded under propositional attitudes verbs, and the resulting sentences are grammatical like in “John is so foolish that he believes that colorless ideas sleep furiously”. If (5) were meaningless, the argument goes, then so would be the more complex sentence containing it. But that is not the case. Another reason would be that we can formulate arguments that show that (5) is false, which we couldn’t do if (5) would be meaningless. For example the truth of “Abstract objects cannot sleep” and of “Ideas are abstract objects” entail the falsity of (5) (Davis, 2003. 262).

Yet, other authors take (5) to be both meaningless and not a member of E (Chomsky 1965). If we follow this later position and if the claim holds for all meaningless expressions (i.e. if an expression is meaningless then it is also not a member of E) then **I** is not a partial but a total function. According to this later position we should treat (5) on a par with “John are nicely”, or “Frightens sincerity John”, that is treat it on a par with other strings of signs that are not members of E (both examples are from Davis 2003).

Normally for the purposes of giving a semantics of a natural language, a theorist will decide whether semantics is a partial or a total function on the set E of expressions. But this is not a decision that I need to take here. I take the issue to be orthogonal with respect to compositionality for two reasons: (a) if compositionality is a constraint on semantics, whether a semantics is partial or total it must satisfy the same constraint and (b) neither partiality nor totality of a semantics guarantee that it is compositional or not. I will leave it open whether a semantics is a total or a partial function on expressions in E. Whatever constraint compositionality sets on a semantics it must be satisfied irrespectively whether the semantics is a partial or a total function, so for the rest of the thesis I assume that anything I say about compositionality applies equally well to semantics that are partial functions and to semantics that are total functions.

Some authors disagree, though (Lahav 1989/2010: 403-404, Recanati 1995: 209). They claim that meaningless expressions like (5) pose a problem for or even falsify compositionality. According to them, compositionality entails that if a complex expression is grammatical (member of E) and has meaningful constituents then the complex expression must be meaningful itself. I think that this is wrong, but I will be in position to discuss it only after I will give a precise formulation of compositionality.

---

Finally, expressions like (5) have correspondent expressions in other languages that mean the same. Intuitively (5) and the Spanish sentence “Ideas incoloras duermen furiosamente” mean the same. If so, both of them are meaningful. (see Magidor 2009). For more arguments in favor of this position see Davis 2003 and Magidor 2009

### 3.3 Assigning Meanings To Expressions: An Abstract Model

Given that semantics assigns meanings to expressions, the semantics of a language can be given by following the same steps as syntax. Syntax specifies a (finite) list of simple expressions and a (finite) list of rules and builds up complex expressions from simple ones with the help of the syntactic rules. Similarly semantics can start by assigning meanings to simple expressions and then build up the meanings of larger ones, from the meaning of their simpler constituents, with the help of semantic rules that determine how meanings of expressions combine.

So we can take our clue from how the syntax was given. Lets start by considering the semantics of a fragment of English (L\*or maybe even a larger fragment of English) that includes the sentence “Mary runs”. One way to start is to do what we did when we specified the syntax. We can start by assuming that the meaning of a sentence is of a certain type **a**, and that the meaning of the name is of a certain type **b**. That is to say that any expression of certain grammatical category S have meanings of type **a** and any expression of category N have meanings of type **b**<sup>7</sup>. Then these two types of meaning constrain what kind of meaning “runs” can have. The meaning of “runs” must combine with the meaning of “Mary” to give the meaning of “Mary runs”. So the meaning of “runs” must be of the appropriate type such that it can combine with the meaning of “Mary” to give the meaning of “Mary runs”. We can say that the meaning of “runs” must be of the type **a/b**, that is of the type that combines with meanings of type **b** to give meanings of type **a**. But, how do meanings combine? Expressions combine through concatenation, but concatenation is not the appropriate operation for combining meanings. One idea is that we construe meanings of expressions as functions. Then if a complex expression that is the result of concatenating one expression with another we can take the meaning of one of its constituents to be a function and the meaning of the other to be an appropriate argument for that function<sup>8</sup>. For example, we can think of the meaning of one constituents of “Mary runs” to be a function

---

<sup>7</sup> In this presentation I somewhat departed from the way in which theorists usually give this kind of syntax. Usually, theorists specify grammatical categories in terms of their semantic import. In fact one of advertised benefits of this approach is that grammatical categories have built-in semantic import.

<sup>8</sup> This idea goes back to Frege (1892/1960). Heim and Kratzer (1998: 13) even call the conjecture that all meaning combination takes place through functional application, Frege's conjecture.

and the meaning of the other constituent to be an appropriate argument for that function, and the meaning of the complex to be the value of that function for that argument. We can take either the meaning of “Mary” (or more generally meanings of type **b**) to be a function and the meaning of “runs” to be in the domain of that function, or the other way around: the meaning of “runs” (or more generally the meanings of type **a/b**) to be a function of the type that has the meaning of “Mary” in its domain. To put it in a nutshell: we can think of meanings as functions and their combination as the operation of functional application. Then we can give a very general rule that determines how the semantics assigns meanings to complex expressions: if  $e_i \wedge e_j$  is a complex expression formed by concatenating  $e_i$  and  $e_j$  and the meaning of  $e_i$  is a function whose domain contains the meaning of  $e_j$  then the meaning of  $e_i \wedge e_j$  is the value of the meaning of  $e_i$  for the meaning of  $e_j$  as an argument:  $\mathbf{I}(e_i \wedge e_j) = \mathbf{I}(e_i) (\mathbf{I}(e_j))$ . This rule establishes how meanings combine. If a semantics **I** assigns meanings to every complex expression according to this rule, then **I** satisfies compositionality: for obviously there is a function such that the meaning of complex expressions is a function of the meaning of their constituents and the syntactic mode of combination, namely the function that applies the meaning of the functor-expression to the meaning of the argument-expression.

In other words, just like in syntax, the idea is to start from a simple base and build successively on that. We can start by assigning certain types of meaning to names (all names have meanings of type **b**) and sentences (all sentences have meanings of type **a**) and then we can construct other types of meanings from the starting point of **a** and **b**. As I just did, “runs” will have a meaning of type **a/b**, that is a meaning that combines through functional application with meanings of type **b** to give meanings of type **a**. (I didn't specify whether **b** is the function and **a/b** the argument or the other way around, because I didn't specify, yet, what **a** and **b** are). All types of meanings that a theorist needs to give an analysis of natural languages can be systematically generated from the basic types **a** and **b**.

Obviously, this description is somewhat general, and when we move from simple toy languages like  $L^*$  to more complex languages further complications will appear. But more importantly this is incomplete in a very important respect. I didn't specify what meanings are, and for this reason I couldn't specify what types of meanings **a** and **b** are. I



will discuss what meanings are in the next section.

## 4. What is Meaning?

### 4.1 Context and Circumstances.

One good starting point to think about meaning is to follow Lewis' advice (1972:194) and claim that meaning *is* what meaning *does*. A rather uncontroversial observation is that what the meaning of declarative sentences do is to determine the conditions under which they are true (or false). Then we can say that meaning of sentences is whatever determines their truth-value given a possible state of affairs.

This basic insight that to know the meaning of a sentence is to know what has to be the case, or how things would have to be in order for the sentence to be true, offers a good starting point for theorizing about meaning. In other words, to know the meaning of a sentence is to know under which conditions the sentence is true. What a semantic theory does, then, is to match sentences in a language with (their) truth conditions and subserviently to match simple expressions with meanings such that given the syntactic rules of the language the meanings of the simples determine the truth conditions of sentences that they are part of. For example, to know the meaning of

(1) "Adrian is a Spanish citizen"

we need to know what has to be the case such that (1) is true. A semantic theory for English, then, produces clauses of the following form:

"Adrian is a Spanish citizen" is true iff Adrian is a Spanish citizen.

in which sentences of the language are paired with their truth conditions. Then the theory will evaluate (1) to be true depending on the way things are. Let's call the complete way things might be a *possible world* and the way things are in fact the *actual world*. One way to cash out the idea that a possible world is needed for truth-evaluation is to model the meaning of a sentence as a function from worlds to truth values. For example, we can

model the meaning of (1) as a function from possible worlds  $w$  to truth values such that it yields truth for a world  $w$  if Adrian is a Spanish citizen in  $w$ . Thus, importantly for subsequent discussion, we say that sentences have a truth value only relative to circumstances of evaluation and that for the purposes of semantic theorizing we can represent their meaning as a function from such circumstances to truth values. Of course, representing circumstances as possible worlds and meaning of sentences as functions from possible worlds to truth values is not the only way to work out the more general idea that the meaning of a sentence yields a truth value only relative to a circumstance of evaluation, but for the purposes of introducing the framework I will follow this line<sup>9</sup>. And although there is no widespread agreement about what goes into circumstances of evaluation thinking about meanings of sentences as functions from circumstances of evaluation to truth value is what might be called the orthodox view<sup>10</sup>. There are different names in the literature for the kind of meaning modeled as a function from circumstances of evaluation to truth values: *semantic content*, *intension*, *proposition*, *what is said*. For the rest of the thesis I will use interchangeably “semantic content”, “proposition” and “intension”.

It is not my aim here to evaluate the virtues or shortcomings of such an approach to meaning. But I should mention that such an account of meaning helps us give an elegant treatment of certain expressions of English. If we take meaning of sentences to be functions that yields different truth values at different possible worlds, then we can treat expressions like “necessary” or “possibly” as intensional operators, namely as functions from such

---

9 For example some theorists prefer to use richer circumstances of evaluation, for example worlds and times, and thus to model meaning as a function from world-time pairs to truth values. I will discuss this issue in more detail in the last section of the chapter.

10 Maybe it is a bit of a stretch to call this the orthodox view. It is one of the two main candidates about what meanings of sentences are. The other candidate represents meanings of sentences as structured propositions. The meaning of (1) is represented as the structured proposition: <Adrian<property of being a spanish citizen>>. This is supplemented with a theory of truth for (structured) propositions. Theorists that prefer structured propositions over truth conditions (Soames 2012 and King 1996) claim that identifying sentence meaning with functions from circumstances to truth-values cannot be a proper theory of meaning because such functions are not fine grained enough so to represent meaning no matter how fine grained circumstances are taken to be.

And I should mention that even some authors (Schaffer 2012) that take meaning of sentences to be truth conditions prefer *not* to model meaning as functions from circumstances to truth values. They claim that meaning of sentences bare truth values absolutely. For example they would say that an actual utterance of sentence (1) in 2012 expresses the meaning that Adrian is a Spanish citizen in 2012 in the actual world. This meaning is true or false depending on the way things are in the actual world in 2012, but it has that truth value absolutely - i.e. relative to any world. We say that such propositions are world-specific: they do not vary their truth value across worlds.

meanings to truth values <sup>11</sup>. What has to be the case, (or how things have to be) such that (1\*) is true?

(1\*)“Necessarily Adrian is a Spanish citizen”

Intuitively, (1\*) is true relative to a possible world  $w$  iff for any possible world  $w^*$  accessible from  $w$ , Adrian is a Spanish citizen in  $w^*$ . Then we can analyze (1\*) as having two immediate constituents: the sentence (1) and the expression “necessarily” and take the meaning of the latter to be a function from the intension of a sentence to truth value: it maps the intension of a sentence to truth iff for any possible world  $w^*$ , the intension is true at  $w^*$ . And it is tempting to believe that other expressions (for example, deontic “ought” and “can”, epistemic “must” and “might”) can be treated in the same manner. For example a theorist can take as a starting point for her account the insight that “ought” and “must” can receive the same treatment as “necessary” only that the accessibility relation is different in their case: a real deontic accessibility for “ought” (e.g.  $w^*$  is deontically accessible from  $w$  iff  $w^*$  does not violate anything that is mandatory in  $w$ ) and epistemic accessibility for “must” (e.g.  $w^*$  is epistemically accessible for an individual in  $w$  iff  $w^*$  is compatible with that individual's knowledge).

More importantly relativising truth to circumstances and modeling meaning as functions from circumstances to truth-values is valuable because it opens the door to a general approach. If a theorist has good reasons to believe that we should extend the set of circumstances beyond possible worlds, so to include, for example, times and locations, then we could give a similar treatment to temporal expressions (e.g. “always”), or locative expressions (e.g. “in Barcelona”). It is a heated debate whether circumstances should be extended beyond possible worlds, and if so what goes into circumstances. I will come back at this in the last section of the chapter.

But when it comes to natural languages this is still unsatisfactory in one important respect. One important feature natural languages is that some of their expressions are context sensitive: whether they are true or false depends on the context of utterance - for example, depends on who the speaker is, or when and where the sentence is uttered. In

---

<sup>11</sup>If one wants to be strict we can say that such operators take a intension as input and yield a intension. For example, “necessarily” takes an intension as input and yields an intension which is true iff the input intension is true relative to all possible worlds.

other words, for some sentences what makes the very same sentence true can vary across occasions of use. For example, the sentence

(2) “I am a Spanish citizen”

does not have truth-conditions absolutely, but only relative to a context of utterance. If uttered by Adrian (2) is false, since Adrian is not a Spanish citizen, but if uttered by Mariano Rajoy it is true. The truth value of (2) depends on who the speaker is, and on whether the speaker is a Spanish citizen and unless we determine the referent of “I” we cannot evaluate (2) for truth. Because of context-sensitivity natural language sentences like (2) have truth-conditions only relative to actual or potential contexts of utterance. Given that “I” varies its reference with the speaker, the state of affairs that make (2) true also vary so a semantic theory will produce the following clause for (2):

“I am a Spanish citizen” is true at a context *C* iff the speaker of *C* is a Spanish citizen at the time of *C*.

Given that we cash in the intuitive idea of truth conditions of sentences by saying that sentences express a content which is (or can be modeled) as a function from circumstances to truth values, then we have to say that some sentences express a content only relative to a context of utterance, and that relative to different contexts of utterance they express different contents. Then, a semantics that assigns contents (functions from circumstances to truth values), will assign different contents to (2) relative to different contexts. For the purposes of this thesis I will work with the following general definition of context sensitivity: a sentence *S* is context sensitive iff it expresses different contents (intensions) across different contexts of utterance, or better put iff utterances of the sentence can be associated with different contents (functions from circumstances to truth values).

If sentences have truth values only relative to contexts of utterance we should give a definition of truth for sentences at a context. We say that for any sentence *S* and context *C*, *S* is true at *C* iff the intension expressed by *S* at *C* yields truth at the world of the context.<sup>12</sup> What is a context? One way to model context is as a quadruple consisting of an

---

<sup>12</sup> Armed with this we can also give a definition of truth of an utterance: we say that an utterance is true iff (the uttered) sentence is true at the world of the context.

agent/speaker, a world, a time and a location,  $\langle a, w, t, l \rangle$  where  $a$  is at  $t$  at  $l$  in  $w$ . Any such a quadruple is a possible context of utterance<sup>13</sup>. I will call the elements of the context, *contextual parameters* and the elements of the circumstance (for now possible worlds) *circumstantial parameters*.

To sum up: when we give a semantics for natural languages, not only that we need to relativise truth (of a sentence) to circumstances of evaluation but we also need to relativize it to context. In order to evaluate (2) for truth value we need both (a) a *context of utterance* to determine who the speaker is and (b) a *circumstance of evaluation* to determine whether the speaker is a Spanish citizen. We say that a sentence is true relative to a context and a circumstance of evaluation. Then, a semantics that assigns intensions will assign them not to sentences but to sentences at contexts (or to sentence-context pairs).

Once we say that sentences have truth-value only relative to a context we can treat other expressions (“here”, “now”, “that”) in a similar fashion with “I”. (Exactly, what expressions can be included in this list is still an open question, to which I will come back at the end of the chapter). In other words, relativising truth of sentences to contexts, allows us to account for context sensitive expressions, just as relativising truth of sentences to circumstances of evaluation allows us to account for modality (and possible other expressions if we can extend the set of circumstances).

This captures another aspect of meaning, different than the one that we model as functions from circumstances to truth values (content). Any competent speaker of English will know under what conditions (2) is true, even if she doesn't know who the speaker is, namely she will know that (2) is true iff the speaker, whoever she is, is a Spanish citizen. Competent speakers of English know this in virtue of knowing that “I” always refers to the person who utters it, and that “I” refers to different individuals if different individuals use

---

<sup>13</sup>We owe this way of representing context to Kaplan (1989). It is not wholly uncontroversial, but it is very suitable for the purposes of introducing the concepts and framework that I'm operating with. As I will proceed, when necessary, I will mention amendments or objections to this way of representing contexts. For example, as a small amendment, some authors argue that we can equally represent contexts as a triple consisting of an agent, world and time, since having an agent in a world at a time determines a unique location. An altogether different way of representing context is as the set of possible worlds determined by commonly accepted presuppositions by the members of a conversation (Stalnaker 1979). Other authors claim that we should distinguish between a narrow notion of context (like the one defined here) and a wider and richer one (the totality of facts relevant in a conversational communication) (K. Bach 2005).

it. This kind of competence is independent of knowing anything about the context of utterance. But this knowledge together with knowledge of context, is sufficient to determine that (2) is true at a context in which Adrian is the speaker iff Adrian is a Spanish citizen. These are the repeatable and conventional aspects of meaning that expressions have independent of context of use and it is what speakers acquire when they learn the meaning of the expression. Following Kaplan I will call this aspect of meaning *character* or *linguistic meaning*. We can model the linguistic meaning of “I” as a function such that for any context of utterance C it yields the speaker or agent of C, and the linguistic meaning of (2) as a function that for any context C yields function that yields truth at those worlds at which the speaker of C is a Spanish citizen.

Finally, the framework allows us to treat some expressions as context-shifting operators. That is, just as the framework allows us to treat certain expressions (e.g. modals) as circumstance-shifting operators, the framework opens up the possibility that some expressions are best treated as context-shifting operators. Although such operators are allowed by the elegance of the framework, whether English (or other natural languages) contain such expressions is not at all a settled issue<sup>14</sup>. Such operators, which Kaplan (1989: 511) calls "monsters", work at the level of linguistic meaning (character) in the same fashion in which circumstance-shifting operators work at the level of content. Consider (2) and the sentence (2\*\*) obtained from (2) by prefixing it with the (rather artificial) context-shifting operator “In some other context”:

(2\*\*) “In some other context I am a Spanish citizen”.

Where (by stipulation) for (2\*\*) to be true in the context of utterance it suffices that some agent of some other context be a Spanish citizen at the time of that context. We can introduce this context-shifting operator (abbreviated as M) in the following way:  $M(S)$  is true at C iff there is a context  $C^*$  (that bears certain relation with C), such that  $S$  is true at  $C^*$ . In the course of evaluation of  $M(S)$  at context C we evaluate S at context  $C^*$ : what the context-shifting operator M does is to look across contexts (that bear a certain relation with C) and yield true if there is a context  $C^*$  such that S is true at  $C^*$ .

---

14 It is undisputed that pure modals can't be monsters. But Perry and Israel (1986) argue that some propositional attitude operators (e.g. "might know") are best treated as monsters.

So we need to distinguish between the linguistic meaning or character of sentences, and their content content relative to a context: for any context of utterance, the linguistic meaning/character determines the content of the sentence at that context, and for any given world the content will determine a truth value. Characters will be modeled as functions from contexts to content and contents will be modeled as functions from circumstances to truth values.

## 4.2. Benefits of Double Indexing

This framework identifies two sources of truth value variation: context and circumstance. If different utterances of the very same sentence have different truth values, it is either because the sentence uttered has different contents at the respective contexts of utterance or because, the content gets evaluated at different circumstances of evaluation. Any theory will decide whether the difference in truth value is because the evaluation of the sentence at different contexts, yields different contents (the sentence contains context-sensitive expressions) or because the content expressed by the sentence gets evaluated at different circumstances (the sentence contains circumstance-sensitive expressions)

I'll briefly mention two other benefits of this framework. We see better the need for double evaluation if we consider two pairs of distinct properties that expressions of natural languages have (in different combinations): rigidity (or circumstance-insensitivity) and non-rigidity (or circumstance sensitivity), on one hand, and context-insensitivity and context-sensitivity, on the other hand<sup>15</sup>.

---

15 For now the only element in the circumstances of evaluation are possible worlds. Then we say that an expression is circumstance-insensitive or rigid iff it denotes the same individual or property relative to any possible world in which the individual or property exists, and denotes nothing in other words. This corresponds to the traditional definition of rigidity which was given in terms of possible worlds. But if we want widen the list of circumstantial parameters (e.g. so to include times, location and maybe other parameters) then we should adjust correspondingly the definition of circumstance-insensitivity. (For example we say that an expression is rigid with respect to time if its character determines a content that is a constant function across times: it yields the same extension relative to any value for the time parameter. Kaplan's example (Kaplan 1989:499) "It is possible that in Pakistan, in five years, only those who are actually here now are envied" shows that "actually", "here", and "now" are rigid with respect to world, time and location: they refer to the world, time and place of the context even if

The difference between them comes clear if we look at (1), (2) and (3).

(1) Adrian is a Spanish citizen

(2) I am a Spanish citizen.

(3) The person speaking here and now is a Spanish citizen.

Consider first (1) and (2). Obviously, the content of (1) does not vary across contexts of utterance: (1) expresses the same content relative to any content of utterance, namely that Adrian is a Spanish citizen (a function that yields true at those possible worlds  $w$  at which Adrian is a Spanish citizen and yields false otherwise). On the other hand, the content of (2) does vary across contexts of utterance if different individuals utter (2). The reason for this is obvious: the proper name in (1) refers to the same individual at every context of utterance, namely to Adrian while the first person pronoun in (2) refers to different individuals at contexts that differ from each other with respect to who is speaking. So proper names and first person pronouns have different properties: proper names do not vary their reference across contexts of utterance while first person pronouns do (we say that proper names are context-insensitive while first person pronouns are context-sensitive).

But the proper name in (1) and the first person pronoun in (2) have one property in common: their reference cannot be shifted by modal operators. For any possible world  $w$  the content of (1) will return truth iff Adrian is a Spanish citizen at  $w$ . In other words, for any possible world it is Adrian (and his properties at that world) that are relevant for the truth evaluation of (1). Which means that the proper name refers to the same individual at any possible world, and thus we say that it is *rigid* or *circumstance insensitive*. This is also true for the first person pronoun in (2): at any context of utterance (2) will express a content that is about the speaker of that context, namely the content that that speaker is a Spanish citizen. We see this better if we embed (2) under a modal operator like in

(2\*) I could have been I am a Spanish citizen.

Sentence (2\*) is true at  $C1: \langle @, \text{Adrian}, 2013, \text{Barcelona} \rangle$  iff there is a possible world  $w$  at which Adrian Briciu is a Spanish citizen. It is the speaker of  $C1$  that is relevant for

---

they are embedded under modal, temporal and locational operators.)



evaluating (2\*), and no other individual. Prefixing “I” with an intensional operator does not change its reference, for once we evaluate "I" relative to a context what is relevant for the truth-evaluation of (2) is still the speaker of that context.

Now compare (2) with (3)

(3) The person speaking here and now is a spanish citizen.

Obviously, (3), just like (2) is context sensitive since different speakers can utter (3) at different times and/or locations, so (3) will express different contents intensions at those different contexts. But there is an easily noticeable difference between (2) and (3). Consider that we evaluate (3) and (2) at the same context C1 (the context consisting of Adrian Briciu, the actual world, 2013, Barcelona). Again the individual relevant for evaluating the intension of (2) at C1 at different worlds is always Adrian Briciu: for any possible world  $w$  the content/intension expressed by (2) be true at  $w$  iff Adrian Briciu is a spanish citizen at  $w$ . The content of (3) at C1 is also true at a world  $w$  at which Adrian Briciu is a spanish citizen. But, as opposed to (2) the content expressed by (3) at C1 is also true at a (non-actual) world  $w'$  at which not Adrian, but a different person is speaking at the time and location of C1 if that person is a spanish citizen at  $w'$ . In other words, if there is a possible world  $w'$ , that differs from the actual world in that not Adrian, but a different individual is speaking in Barcelona on the 1<sup>st</sup> of October 2013, the intension of (2) at C1 is false at  $w'$ , while the intension of (3) is true at  $w'$ . This too can be seen better if we embedd (2) and (3) under modal operators, like in

(3\*) The person that speaks here and now could have been a spanish citizen

and

(2\*) I could have been a spanish citizen”

and we evaluate both at C1. The content of (3\*) at C1 is true iff there is a possible world  $w$ , such that the individual that utters (3) in Barcelona at 2013 is a spanish citizen in  $w$ . That individual need not be Adrian Briciu. On the other hand (2\*) is true at C1 iff there is a possible world in  $w$  such that Adrian (and no other individual) is a a spanish citizen in  $w$ .

So, this framework allows us to distinguish between different kind of properties that expressions of natural language have: context-sensitivity (and context-insensitivity), on one

hand, and circumstance-sensitivity (and circumstance-insensitivity) on the other hand. Context-sensitivity and circumstance-sensitivity are properties of different aspects of meaning. Context-sensitivity is a property of linguistic meaning and circumstance sensitivity is a property of content. For the purposes of semantic theorising we can model the character of *context-insensitive* expressions as constant functions that yield the same content relative to any context and the character of *context-sensitive* expressions as a non-constant function that yields the varying intensions across contexts of utterance. Similarly, the content of *circumstance-insensitive* expressions can be modeled as constant functions that yield the same extension at any circumstance of evaluation, and those of *circumstance-sensitive* expressions as non-constant functions of this type.

Another benefit of double evaluation is that it helps explain the difference between *necessity* (truth at all possible worlds) and *analyticity* (truth at all contexts of utterance)<sup>16</sup>. It can explain why some sentences are analytic but express contingent contents. For example, (3) is true at every context, since for any context C the speaker of C is at the time and location of C<sup>17</sup>.

(4) "I am here now"

But given any context (4) does not express a necessary truth. In other words, this framework allows us to explain both the intuition that (4) cannot be false since it is true in virtue of its meaning, and the intuition that

(4\*) "Necessarily I am here now"

is false since, obviously, the speaker could have been somewhere else at the time of utterance. It explains the first intuition by assigning to (4) truth relative to every context - or rather, by assigning to every pair consisting of (4) and a context, an intension that is true at the world of the context. This explains its analyticity.

It explains the intuition that (4) doesn't express a necessary truth in a simple way. Whenever a sentence is in the scope of the necessity operator, its intension is evaluated relative to all possible worlds. If the evaluation yields truth relative to all possible worlds,

---

16 What I call analyticity Kaplan calls logical truth.

17 Alternatively, we can say that every utterance of "I am here now" is true, if we define truth for utterance as: an utterance is true iff (the uttered) sentence is true at the world of the context.

the sentence expresses a necessary truth. But, for example, relative to a context <Adrian, @, Barcelona, 2013> (4) expresses an intension that does not yield truth at every possible world since there is a possible world in which Adrian Briciu is not in Barcelona in 2013. In other words, prefixing (4) with “necessarily” we obtain a sentence (4\*) that expresses a falsehood, and therefore (4) does not express a necessary truth.

### 4.3. Predictions of Truth for Utterances

An important question is how we go from assigning meanings/contents to sentences relative to contexts to predicting the truth values of utterances of natural languages (what semantic theories aim at). Utterances are real events in which what get uttered are expressions of natural languages (e.g. English). On the other hand semantics (as defined in the first part of the chapter) assigns meanings not to English expressions, but to the expressions of a disambiguated language (what I called *the model-language*). For the purposes of semantic theorizing we can represent utterance events as pairs consisting of an expression and a context of utterance, where the expression in question is an expression of the disambiguated language and the context is a theoretical construct meant to represent a real speech situation.

Our semantic theory will make predictions about these sentence-context pairs. But given that the aim of semantic theories is to predict the truth values of utterances of natural language sentences and that semantic theories have empirical import only inasmuch as they make such predictions, whether a semantic theory is successful or not depends the representation relation it establishes between sentence context-pairs and concrete utterance events. Only given certain assumptions about the relation between concrete utterance events and their theoretical representation as sentence-context pair, can a theorist claim that the predictions of the semantic theory are predictions about concrete utterance events.

Two things are needed in this direction: (a) to specify how contexts as theoretical constructs (i.e. qua tuple of parameters) can represent real speech situations and (b) to specify truth for sentence-context pairs (which I already did). The relation between context

of utterance qua real speech situation and context qua theoretical object can be thought of in the following way: given an utterance event  $u$  the agent of the context (qua theoretical object) is the person who is producing  $u$ , and the time, location and world of the context are the time, location and world respectively at which  $u$  takes place<sup>18</sup>. What relies on some further considerations from the part of the theorist is which context *qua* theoretical object and which sentence of the model language should be the constituent of a sentence-context pair that represents a particular utterance event.

The second element needed to tell how predictions about sentence-context pairs give predictions about concrete utterance events is a definition of truth for sentence-context pairs. This was already provided: we say that  $\langle S, C \rangle$  is true, or that sentence  $S$  is true at a context  $C$  iff content assigned to  $S$  at  $C$  yields true at the world and time of the context.

And now we can specify the relation between a concrete utterance event  $u$  and a theoretical construct (a sentence context pair  $\langle S, C \rangle$ ) meant to represent it. We say that a sentence context pair  $\langle S, C \rangle$  represents an a real utterance event  $u$  iff the intension assigned to  $S$  at  $C$  is true iff  $u$  is true. Given that we have pre-theoretical access to intuitions about truth values of utterance  $u$  we can, then, test the predictions made by the semantic theories. For example, a theory makes wrong predictions if on one hand speakers have the intuition that an utterance of (2) by Barack Obama is false and on the other hand the theory assigns to a sentence-context pair of which its constituents are (2) and  $C_2$  (where  $C_2$  is such that  $a_{c_2} =$  Barack Obama, and  $w_{c_2} =$  the actual world) a content that yields true at the actual world. This looks intuitively correct: an actual utterance of (2) by Barack Obama is true iff Barack Obama is a spanish citizen in the actual world, and since he is not, such an utterance is false. When it comes to giving correct predictions of truth values, we are especially interested in evaluating the content of a sentence-context pair at the world of the context. That's because any utterance event takes place at a world, and we judge an utterance as true or false depending on how things are at that world.

So, context plays a plays a *double role*. First, it provides the interpretation for context sensitive expressions (e.g. "I", "now") that is, it determines the intension of

---

18 For a discussion see Predelli, 2005: 23-25

sentences like (2). Secondly, it determines what circumstances of evaluation are relevant when it comes to generating predictions about truth-values of utterances.

#### 4.4 Choice of Semantic Value

What kind of entities can a semantics assign in order to generate truth value predictions for a language that contains context-sensitive expressions? (Since such entities are the value of the semantic function (i.e. the interpretation function) I will follow the custom of calling them semantic values.). This general framework allows several choices of semantic values, so I'll distinguish between different semantics', that is between different types of functions that assign different types of entities. For this reason I will use subscripts to differentiate between them.

##### 4.4.1. First Option: Extensions

One option is that a semantics assigns directly truth values. Since a sentence is true relative to a context of utterance and a circumstance, such a semantics assigns to each sentence-context-circumstance triple a truth value. And to each sub-sentential expressions it assigns relative to a context and a circumstance an extension. For example, to “I” in (2) it assigns relative to C, at which Adrian Briciu is the speaker, and possible world  $w$  the individual Adrian Briciu, and to “is a Spanish citizen” it assigns relative to C and  $w$  the set of individuals that are spanish citizens in  $w$ . And for any context C and world  $w$  it assigns truth to (2) “I am a Spanish citizen” at C and  $w$  iff Adrian Briciu is a member of the set of individuals that are Spanish citizens in  $w$ . In other words, such a semantics is defined over the Cartesian product of E (the set of expression) and C (the set of contexts) and Circ, (the set of all possible worlds) and with values in the set of extensions  $\{M_{EXT}\}$ :

$$I_{EXT}: EXCX \text{Circ} \rightarrow \{M_{EXT}\}$$

##### 4.2.2. Second Option: Content (Intension)

The second option is semantics that assigns as values to sentence-context pairs intensions: functions from circumstances to extensions (truth values). And more generally, such a semantics assigns to every expression-context pair a function from circumstances to

extensions: to singular terms functions from circumstances to individuals (or objects) and to predicates functions from circumstances to sets of individuals. This semantics is defined over the Cartesian product of E and C and has value on the set  $\{M_{\text{int}}\}$  where the latter is the set of functions from circumstances to extensions

$$\mathbf{I}_{\text{INT}}: \text{ExC} \rightarrow \{M_{\text{EXT}}\}.$$

#### 4.2.3. Third Option: Linguistic Meaning (Character)

As a third option is a semantics that assigns as meanings characters: functions from contexts to intensions/content. A semantics that assigns characters to expressions is a semantics from expression in E with values in  $\{M_{\text{LM}}\}$  where  $\{M_{\text{LM}}\}$  is the set of functions from contexts to functions from circumstances of evaluation to truth values:  $\mathbf{I}_{\text{LM}} : E \rightarrow \{M_{\text{LM}}\}$ .

#### 4.2.4. The Relation Between The Three Choices

The following relations hold between these three types of semantics'. The first semantics is a function from expression-context-circumstances triples to extensions

$$\mathbf{I}_{\text{EXT}}: (\text{EXCXCirc}) \rightarrow M_{\text{EXT}}$$

Let  $S$  be a sentence such that  $S \in E$  and  $C \in C$  and  $w \in \text{Circ}$ , then  $\mathbf{I}_{\text{EXT}}(S, C, w) = 1$  (true) iff  $S$  is true at  $C$  and  $w$

There is a well-known logical procedure, schönfinkelization (or currying), by which an  $n$ -place function is reduced to a one placed function. For example, we can reduce the two place function  $f(x,y)$  to the one place function  $f(x)(y)$ . We say that  $f(x)(y)$  is the schonfinkelized version of the two place function  $f(x,y)$ .

By schonfinkelizing  $\mathbf{I}_{\text{EXT}}$  we get an associated semantics defined over expression-context pairs and has values in the set of intensions:

$$\mathbf{I}_{\text{INT}}: (\text{EXC}) \rightarrow M_{\text{INT}}$$

Let  $S, C$  and  $w$  be as above then:

$$\mathbf{I}_{\text{INT}}(S, C) = \text{a function from circumstances to truth values such for any } w, \mathbf{I}_{\text{INT}}(S, C)(w) = 1 \text{ iff } \mathbf{I}_{\text{EXT}}(S, C, w) = 1$$

Such a semantics assigns to each sentence context pair  $\langle S, C \rangle$  a function defined over the

set of circumstances and with values in the set of truth values, such that for any circumstance  $w$ , the function yields true iff  $\mathbf{I}_{\text{EXT}}(S, C, w) = 1$ .

Finally, by schonfinkelizing  $\mathbf{I}_{\text{INT}}$  we get a semantics is defined over expressions and has values in the set of characters:

$$\mathbf{I}_{\text{LM}}(E) \rightarrow M_{\text{LM}}$$

Again let  $S, C$  and  $w$  be as above then

$\mathbf{I}_{\text{LM}}(S)$  is a function from contexts to intensions such that for any context  $C$ ,  $\mathbf{I}_{\text{LM}}(S)(C)$  yields a function  $\mathbf{I}_{\text{INT}}(S, C)$  such that  $\mathbf{I}_{\text{INT}}(S, C)(w) = 1$  iff  $\mathbf{I}_{\text{EXT}}(S, C, w) = 1$

It is easy to see that for any  $S$  and any  $C$ :  $\mathbf{I}_{\text{LM}}(S)(C)(w) = \mathbf{I}_{\text{INT}}(S, C)(w) = \mathbf{I}_{\text{EXT}}(S, C, w)$ .

#### 4.5 Extending the Framework: Choices in Modeling Content

As I already said given that sentences have truth-value only relative to a context and a circumstance of evaluation it is clear that divergence in the truth-values of different utterances of the very same sentence can be due to either the fact that (a) the semantic content of the two utterances differ (the uttered sentence expresses different contents at the two contexts) or (b) the semantic content of the two utterances is the same, but it is evaluated at different circumstances. For example two utterances of

(5) "I am hungry"

by Adrian made at different contexts  $C_1$  and  $C_2$  such that the contexts differ only with respect to time ( $t_{c_1} \neq t_{c_2}$ ) can have different truth values. Given that, so far, I take circumstances of evaluation to consist in possible worlds only, we can give the following account for the differences in the truth-values of (5). Since we model semantic content as a function from possible worlds to truth values. and the semantic content of the two utterances is evaluated relative to the same possible world (the actual world) the difference in truth value must be due to the difference content. The utterance of (5) at  $C_1$  expresses the

content that Adrian is hungry at  $t_1$ , which is (or determines) a function from possible worlds to truth values such that it maps a world  $w$  into truth iff John is hungry at  $t_1$  in  $w$ . The utterance of (5) at  $C_2$  expresses the content that Adrian is hungry at  $t_2$ , which is (or determines) a function from possible worlds to truth values, such that it maps a world  $w$  to truth iff Adrian is hungry at  $t_2$  in  $w$ . That is, the explanation for the differences in truth-values of utterances of (5) is that they express different contents.

But the above general framework allows for another way to account for the differences in truth-values of utterances of (5). We can take circumstances to be richer than just possible worlds, for example we can take circumstances to consist of worlds and times. We could, in other words, extend the list of circumstantial parameters so to include other elements than possible worlds.

Then, the content expressed by a sentence is evaluated relative to a possible world and a time and it will be represented as a function from world-time pairs to extensions. Thus we can take the semantic content of (5) in  $C_1$  and in  $C_2$  to be the same but claim that the circumstances of evaluation of are different. Both utterances express the content that Adrian is hungry, which is (or determines) a function such that when evaluated at the different circumstances  $\langle w, t_1 \rangle$  and  $\langle w, t_2 \rangle$  yields different truth values if Adrian is hungry at one world-time pair but not at the other. Then we explain the difference in truth-values of utterances of (5) by claiming that they express the same content, but the circumstances of evaluation are different.

Of course, a theorist must offer an argument for extending the set of circumstances from possible worlds to possible worlds and times (the set of circumstances now is the cartesian product of the set of worlds and the set of time-moments). One such argument, due to Kaplan (1989) is the following (briefly presented). Consider sentence

(6) "I am hungry now"

If uttered at context  $C_1$ , (6) is true iff Adrian is hungry at the time of  $C_1$ , and if uttered at context  $C_2$ , it is true iff Adrian is hungry at the time of  $C_2$ . So just like (5), it also varies its truth-value with respect to time. But we should treat (5) and (6) differently. The difference between them lies in that (5) can be embedded under temporal operators ("expressions like "sometimes", "every time", "every Sunday", "some days", etc) while (6) can't be embedded



under them:

(5\*) “Sometimes I'm hungry”

is meaningful, (it is true at C1 iff there is a time  $t$  such that Adrian is hungry at  $t$ ) while

(6\*) “Sometimes I'm hungry now”

is not. A plausible explanation for this is the following: expressions like “sometimes”, “every time” (and others) work like temporal operators in the same manner in which “possible” or “necessary” work like a modal operators, they shift the time at which we evaluate the sentence embedded under them. To be able to evaluate the content of a sentence relative to some time moment, then the content itself must be temporal specific, otherwise the evaluation is vacuous. The content of (6) at C1 seems to be temporal specific thanks to the value of “now” at the context C1, and that is why it can't be embedded under “sometimes”. On the other hand, the content of (5) can be embedded under temporal operators and that suggests that it is temporally unspecific. Or in other words, if the job of the temporal operator is to make specific some temporal unspecificity in the sentence that they embed there is a job for them to do only if the content is temporal unspecific. Since (6) is already temporal specific thanks to “now”, there is no job for them to do, and (6) cannot be embedded under such operators. On the other hand (5) is not temporal specific and that is why it can be embedded. But then, for (5) not to be temporal specific, it must be that contents are not functions from possible worlds to extensions, for these functions are temporal specific, but functions from world-time pairs to extensions.

This argument can work as a template for other arguments in favor of extending the set of circumstances with new parameters. For what I said above might not be restricted to worlds and times, but it concerns anything that is relevant in determining the truth values of an uttered sentence. The same recipe applies to other elements relevant in the determination of truth-values: if  $x$  is a determinant of truth value, then  $x$  “is *either* given as an ingredient of content *or* as an aspect of the circumstance of evaluation” (Recanati: 2007, 33-34). And, of course, some theorists want to take circumstances to be even richer than worlds and times. How we decide to model content and how rich or sparse we take circumstances to be are different faces of the same coin. Consider another example and possible extension. The sentence

(7) “It is raining”

can be true if uttered in Barcelona on a rainy day, but false if uttered on the same but sunny day in Paris. It seems that where the utterance (7) takes place is relevant for its truth value. We can account for this difference in truth value, either by claiming that the two utterances express different contents or, alternatively, we can say that the two utterances of “It is raining” express the same content but the relevant circumstances of evaluation are different, and thus determine different truth values. If we take the first option we’ll say that when uttered in Barcelona (7) expresses the content that it is raining in Barcelona which is, or determines, a function from world-time pairs to truth values such that it yields truth for  $\langle w, t \rangle$  iff it is raining in Barcelona at  $t$  in  $w$ . When uttered in Paris (7) expresses the content that it is raining in Paris, which is, or determines, a function from world-time pairs to truth values such that it yields truth for  $\langle w, t \rangle$  iff it is raining in Paris at  $t$  in  $w$ . If so, such a type of content does not vary in truth values across locations, but only across worlds and times. We’ll say that content is location-specific. On the other hand, if we take the second option, the two utterances of (7) express the same content, namely that it is raining, which is, or determines, a function from world-time-location triples to truth values such that it yields truth for  $\langle w, t, l \rangle$  iff it is raining at  $t$  at  $l$  in  $w$ . Such content varies its truth values not only across worlds and times but also across locations (what is called in the literature a *relative proposition*)<sup>19</sup>.

This is not the place to go over the arguments offered in favor of one or another way of modeling semantic content. Nor will I take sides in this debate. It is important to mention though that there are several ways in which one can define semantic content, since this is important for a precise definition of compositionality. It is also important for another reason related with compositionality. There is a general strategy that we can use to save compositionality in the face of looks like cases of failure of compositionality for semantic content: we choose a different type of semantic content. For example, if a semantics of English (or a fragment of it) assigns temporal propositions (functions from world-time pairs to truth values) and assigns different contents to (7) at the Barcelona-context and at the

---

19 And of course some theorists take circumstances to be even richer than world-time-location triples. Recent debates about how to model semantic content to account for the phenomenon of context sensitivity have spawned a considerable literature on this subject. See Kölbel (2008) for an overview.

Paris-context, but it assigns constant contents to its constituents across the two contexts, then the semantics fails to be compositional. But failure of compositionality can be avoided by having a semantics that assigns relative propositions (functions from world-time-location triples to truth values), that assigns constant contents to (7) and to its constituents across contexts. (I will come back on this in the next chapter).

Choice of how to model content, and more generally choice of semantic value, will be guided not just by considerations regarding predictions of truth values but also by a variety of other considerations and desiderata. Predicting correct truth values to utterances is the minimal requirement that we can place on a semantic theory. Some authors might want semantic theories to do more than just model linguistic competence. For example, a theorist might expect semantic theories to play a role in explaining propositional attitudes states, or to play a role in explaining successful linguistic communication, or a role in explaining how speakers use sentences to make assertions, or a role in explaining the intuition that on some topics people can disagree without any of them being at fault, and so on. For example if a theorist believes that a semantic theory should play a role in explaining how speakers use sentences to make assertions, then she might be inclined to choose as semantic values those type of entities that can also work as objects of assertions.

But I should stress that in this thesis I do not place additional explanatory burdens on semantic theories. For example, in this thesis I will work under the assumption that the semantic value of a sentence at a context and the asserted content of the corresponding utterance can be distinct and that they play different theoretical roles. In other words I don't take, as some theorists do (e.g. Kaplan 1989), the semantic value assigned by semantics to a sentence *S* at a context *C* to be identical with the content asserted by uttering *S* at *C*.

## **5. Summing up**

Now that I specified what meaning is, in fact what choices of semantic values a theorist has, it is time to tie some loose ends. It is time to go back and see how semantics assigns meanings in a compositional manner. I said that one starting point is to assign certain type of meaning to names (all names have meanings of type **b**) and certain type of

meaning sentences (all sentences have meanings of type **a**) and then we can construct other types of meanings from the starting point of **a** and **b**, but I didn't specified what **a** and **b** are. Now I'm in position to do that.

Consider a semantics that assigns extensions as meanings. Then, to a name it will assign an individual as semantic value (extensions of type **a** are individuals) and to a sentence it will assign a truth-value (extensions of type **b** are truth-values). Then the extension of type **a/b** is a function that maps individuals into truth values. Take the sentence “Mary runs”. To “Mary” it will assign as semantic value the individual Mary. To “Mary runs” it will assign truth iff Mary runs. And to “runs” it will assign as semantic value a function from individuals to truth values, such that it will map every individual  $x$  into truth if  $x$  runs and it will map into falsehood otherwise.

Obviously, a semantics that assigns intensions will work very much in the same way. To a name it will assign an individual as semantic value, or equivalently a constant function from circumstances to that individual. So semantic values of type **a** are individuals – or constant functions from circumstances to individuals. To a sentence it will assign as semantic value an intension: a function from the set of circumstances to truth-value. So semantic values of type **b** are such functions. Then the semantic value of type **a/b** is a function from individuals to functions from circumstances to truth values. So, “Mary runs” is assigned an intension: a function from circumstances to truth values, that yields truth at those circumstances at which Mary runs. The name “Mary” is assigned as semantic value the individual Mary (or a constant function from the set of circumstances to the individual Mary). And “runs” is assigned as semantic value the function that maps every individual  $x$  to a function that maps circumstances into truth-values such that it will map a circumstance  $Circ$  into truth iff  $x$  runs at  $Circ$ <sup>20</sup>.

Both in the extensional and in the intensional semantics sketched above the assignment of semantic value to “Mary runs” is compositional: the semantic value of “Mary runs” is function of the semantic values of its constituents and the way they are combined. This offers a blueprint or a sketch of how a compositional semantics for (a larger

---

<sup>20</sup> And, rather obviously, a semantics that assigns characters works in a similar fashion, with the addition of an extra “epicycle” in the assignment of semantic values.

fragment of) English might look like if we extend this model<sup>21</sup>.

Finally, formal semantic theories are taken to model competence with a language by means of generating truth value predictions for sentences of a model language that stands in a certain relation (i.e the actual-language relation) with the language used by a group of speakers whose competence is modeled. The model language is an abstract object: its properties are either stipulated or derived from stipulations and they are, in a sense, independent of empirical facts. The statements and predictions of the semantic theory are about this abstract object. What is empirically determined, through the actual-language relation, is whether the language as described by the semantic theory is used by a certain group of people. If the language as described by the theory *is* the actual language used by some population, then generating systematic truth-value predictions by the semantic theory achieves two things: (a) it gives a characterization of the semantic and logical relations that hold between the expressions of the language used by the population (ambiguity, synonymy, logical truth, entailment, rigidity, etc) and (b) it models the linguistic competence of members of that population.

Predictions of truth-values gives us a good characterization of the semantic properties of a language since from such predictions we can derive predictions about synonymy, entailment and other properties and relations. For example, if a semantics is such that for any sentence  $S$  that contains the expression  $e_i$ , and any sentence  $S'$  obtained by replacing  $e_j$  for  $e_i$  in  $S$ , and for any context  $C$ , it predicts that the truth value of  $S$  at  $C$  is identical with the truth value of  $S'$  at  $C$  then the semantics establishes that  $e_i$  and  $e_j$  are synonymous (namely that they have the same conventional meaning). Or to put it somehow closer to the framework used here: a semantics that assigns semantic values to expression-context pairs, and that predicts the same truth value for the pairs  $\langle S, C \rangle$  and  $\langle S', C \rangle$  for any  $S$  and  $S'$  and any  $C$ , where  $S'$  obtained by replacing  $e_j$  for  $e_i$  in  $S$ , also predicts that  $e_i$  and  $e_j$  are synonymous.

---

21 In this essay I take the simplest semantic units to be words and not lexemes and morphemes. So in what follows I will treat expressions like "runs" and "running" as semantically simple. For example, I will not consider the root-morpheme RUN and the derivational morpheme -ING to be the basic meaningful units of the language. I'm not concerned with whether RUN and -ING have meaning and with *lexical compositionality*, that is with whether and how the meaning of "running" is determined from the meaning of RUN and -ING. I'll take no stand on whether the meaning of syntactically simple but morphologically complex expressions is determined from the meanings of their lexemes and morphemes.

When it comes to linguistic competence we start from the insight that knowledge of meaning is, or involves, knowledge of the conditions under which a sentence is true, and we say that a user is competent with a language only if she can specify for any arbitrary meaningful sentence *S* of that language how the world has to be in order for *S* to be true. That is, one counts as competent with that language only if given full knowledge of the relevant worldly facts she can evaluate correctly the truth-value any actual or possible utterance of *S*. Semantic theories model this competence by pairing sentences with their linguistic meaning which, in its turn, determines under which conditions a sentence is true or false (or alternatively by pairing sentence-context pairs with truth-conditions).



## CHAPTER 2: Principles of Compositionality

### 1. Introduction

I can, now, take the first steps towards a precise formulation of compositionality<sup>22</sup>. Compositionality introduces a condition on semantics: it constrains the assignment of meaning to complex expressions. So I'll talk about a semantics that obeys the principle of compositionality as a semantics that has the property of being compositional: compositionality is a formal property of the semantics. A fairly common and traditional formulation of the principle of compositionality is the following:

*The meaning of a complex expression is a function of the meanings of its constituents and the way they are combined.*

This formulation has strong intuitive appeal but, as it stands it can be read in several ways and thus needs to be made more precise.

Many texts that discuss the principle of compositionality<sup>23</sup> consider different formulations of the principle found in the literature, formulations of the type:

(a) “a language is compositional if the meaning of each of its complex expressions is *derived* from the meanings of its simple constituents” (Dever 2005: 633) or

(b) “a sentence of [a language] L is compositional if and only if a (canonical) representation of its linguistic structure encodes all the information that a speaker/hearer of L requires in order to understand it” (Fodor and Lepore 2004: 77).

I take these formulations to be either rough approximations, or sketches, of the principle given above (this is the case of (a)), or to formulate other principles (this is the case of (b)). Since it is not my purpose here to discuss various rough sketches of the principle, or to

---

22 Some of the issues that I discuss in this chapter are also discussed in Westerståhl 2009, Pagin and Westerståhl 2010a,b, Pagin and Pelletier 2007, Szabó 2013.

23 e.g. Dever 2005, Jansen 1983 and 1997, Pelletier 1994, Szabó 2000 and 2013



discuss other principles and how they relate with compositionality I will settle for the above formulation of the principle of compositionality.

The problem under consideration in this chapter, though, is how to give a formal specification to the above formulation of the principle of compositionality. The principle can be read in different ways, and thus it is amenable to different (non-equivalent) formal specifications. This is what I will do in the first section of this chapter.

## 2. Formulating Compositionality

As I already mentioned the general formulation of compositionality, namely that there is a function from the meanings of constituents and syntactic structure to the meaning of the complex can be understood in different ways. In the following section I'll try to give a precise mathematical formulation for the principle of compositionality. As I'll point out in the fifth chapter this precise formulation is, still, somewhat unconstrained and can be strengthened by placing various limiting conditions on compositionality.<sup>24</sup>

A more precise formulation of the principle is needed but I must proceed with care so not to give a formulation so broad that it is trivially true nor to give a formulation so strong that it is obviously false.

For example, one formulation is the following (call it AAE compositionality<sup>25</sup>):

**I** is compositional iff for any syntactic rule  $\alpha$  and any arbitrary expressions  $e_i, e_j \in E$  if  $\alpha(e_i, e_j)$  is meaningful there is a function  $f$  such that  $\mathbf{I}(\alpha(e_i, e_j)) = f(\alpha, \mathbf{I}(e_i), \mathbf{I}(e_j))$

According to this formulation for any pair consisting of a syntactic operation  $\alpha$  and

---

24 For the purpose of simplifying the exposition (following Pagin and Westerstahl (2010a)) I will call a function that takes a pair of meanings (or alternatively, a pair of meanings and syntactic rule) and yields a meaning a *composition function*

25 I use capital letters to highlight the order of the quantifiers (universal (A) and existential (E)) in the different ways of making the principle of compositionality precise. Variables  $e_i, e_j, en$  etc. vary over elements of E (the set of expressions of a language) and variables  $\alpha, \beta$ , etc. vary over elements of F (the set of syntactic rules of a language)

any arbitrary expressions  $e_i, e_j$  if  $\alpha(e_i, e_j)$  is meaningful there is a composition function that determines the meaning of  $\alpha(e_i, e_j)$ . This principle only requires that for any pair  $\langle \alpha \langle e_i, e_j \rangle \rangle$  if the expression  $\alpha(e_i, e_j)$  obtained by combining the two expressions according to the syntactic rule is meaningful, then there is a function from the meanings of its constituents and its syntactic structure to the meaning of the complex. In fact, since this principle allows that *for each complex expression* there is a function that determines its meaning in terms of the meanings of its constituents and its syntactic structure, then *any semantics satisfies it*. In other words, this is a trivial or vacuous principle, and is absolutely useless for the purposes of semantic theorizing.

Another possible precification of the principle of compositionality is the following (call it EAA compositionality):

**I** is compositional iff there is a function  $f$  such that for any syntactic rule  $\alpha$  and any arbitrary expressions  $e_i, e_j \in E$  if  $\alpha(e_i, e_j)$  is meaningful then  $\mathbf{I}(\alpha(e_i, e_j)) = f(\mathbf{I}(e_i), \mathbf{I}(e_j))$

A semantics fails to satisfy this principle if it assigns different meanings to two expressions  $\alpha(e_i, e_j)$  and  $\beta(e_i, e_j)$  that have identical constituents but different syntactic structures. For example although they have the same constituents the expressions “the fat average Romanian” and “the average fat Romanian” have different meanings because they have different syntactic structures. It is clear that this principle is not suitable for semantics of natural languages, since natural language obviously don't satisfy it.

The formulation that we're looking for must be the following (call it AEA compositionality):

**I** is compositional iff for any syntactic rule  $\alpha$  there is a function  $f$  such that for any arbitrary expressions  $e_i, e_j \in E$  if  $\alpha(e_i, e_j)$  is meaningful then:  $\mathbf{I}(\alpha(e_i, e_j)) = f(\alpha, \mathbf{I}(e_i), \mathbf{I}(e_j))$

Here the existential quantifier that introduces the composition function is in the scope of the universal quantifier that ranges over syntactic rules. According to AEA for any syntactic rule  $\alpha$  there is a function  $f$  such that for any expression built by that rule, if the expression is meaningful, its meaning is an  $f$  function of the meanings of its constituents and the way they are syntactically combined. A semantics fails to be AAE compositional if it assigns different meanings to two complex expressions with identical syntactic structure

and pairwise constituents with identical meanings. That is, a semantics fails to satisfy the AEA principle if  $\mathbf{I}(e_j) = \mathbf{I}(e_n)$  and  $\mathbf{I}(\alpha(e_i, e_j)) \neq \mathbf{I}(\alpha(e_i, e_n))$ .

The principle that best captures the intuitive idea behind the principle of compositionality and that looks promising for the purposes of giving a semantics of natural languages is AEA (and its orthographical variants). In what follows, this is the reading of compositionality that I will use.

Now, another question is whether the explanatory purposes of compositionality can be served by a mere functional relation between the meaning of complexes and the meaning of their constituents plus syntactic structure. Most theorists take compositionality to be an *explanatory principle*, that the semantics of a language satisfies compositionality is taken to explain certain properties of the language and, certain features of linguistic competence with that language. In other words, given that compositionality is taken to explain certain linguistic properties, can the definition of compositionality given above its expected explanatory role? Some theorists might doubt that the mere fact that the meaning of complex expressions is a function of the meaning of their parts and syntactic structure can explain those linguistic properties that compositionality is expected to explain.

For this reason some authors have argued that in order for compositionality to fulfill its explanatory purposes we need to place additional constraints or restrictions on the meanings assigned (i.e. what kind of entities can be assigned as meanings) or on what functions are acceptable composition functions (i.e. what operations that combine meanings are allowed). For example, we should require, at least, that the composition function employed by a semantics be computable. Other authors with really demanding expectations about what compositionality should explain have argued that even stronger constraints need to be placed. For example, Szabó (2000) argues that we should require that all possible human languages employ the same composition function. That is to require that any two possible human languages that share the same syntactic rule associate with it the same composition rule (i.e. the same way of combining meanings).

This is not the place to go over such arguments nor to discuss what additional constraints are needed to strengthen the principle of compositionality so that it can serve its explanatory purposes. I'll come back to this in the last chapter and discuss the issue at

length.

### 3. Compositionality and Semantic Values

I argued above that of the different ways of specifying the determination relation stated by the principle of compositionality, the best is the following: for any syntactic rule  $\alpha$  there is function  $f$  such that for any two expressions  $e_i, e_j$ , if  $\alpha(e_i, e_j)$  is meaningful then  $I(\alpha(e_i, e_j)) = f(\alpha, I(e_i), I(e_j))$ .

Obviously such a principle can be straightforwardly applied to a semantics that assigns meaning to expressions independently of context<sup>26</sup>. For semantics' that assign meanings to expressions relative to context we need to reformulate the principle of compositionality. What is important for each of the following formulations of compositionality is that they should capture the general idea that the meaning of a complex expression is determined by the meanings of its constituents and of its syntactic structure and the precise determination relation specified by the AEA principle.

As I emphasized in the previous chapter a theorist has several options to choose from: she can take meanings to be extensions (assigned to expressions relative to a context and a circumstance of evaluation), or to be functions from circumstances to extensions (assigned relative to expression at a context) or to be functions from contexts to functions from circumstances to extensions (assigned to expressions).

---

26 Sometimes this point is confused with another one. Some authors say that this principle holds only for a semantics that assigns context-insensitive meanings (and by extension only to semantics of languages that do not contain context-dependent expressions) and that is not suitable for a semantics that accounts for context-dependency. I think this is wrong. The reason is because a semantics that assigns meanings independently of contexts can account for context sensitivity by simply assigning meanings that are non-constant functions that take context as argument (e.g. by assigning characters that are not constant functions).

### 3.1 Compositionality for Linguistic Meaning (Character Compositionality)<sup>27</sup>

I will start by presenting a formulation of compositionality for the type of semantics described the last section of the previous chapter, namely a semantics that assigns linguistic meaning (character), modeled as functions from contexts to intensions, as semantic values.

According to *the principle of compositionality for characters* the character of a complex expression is a function of the characters of its constituents and of its syntactic structure.

A more precise formulation of the principle is the following: a semantics  $\mathbf{I}^*$  is character compositional iff for any syntactic rule  $\alpha$  there is function  $f$  such that for any two expressions  $e_i, e_j$  if  $\alpha(e_i, e_j)$  is meaningful then  $\mathbf{I}^*(\alpha(e_i, e_j)) = f(\alpha, \mathbf{I}^*(e_i), \mathbf{I}^*(e_j))$ .

A semantics *fails to be character compositional* if for some expressions  $e_i, e_j, e_n$ , and syntactic rule  $\alpha$ , (a) and (b) hold: (a)  $\mathbf{I}^*(e_j) = \mathbf{I}^*(e_n)$  and (b)  $\mathbf{I}^*(\alpha(e_i, e_j)) \neq \mathbf{I}^*(\alpha(e_i, e_n))$ .

I will call the above statement *the failure condition for character compositionality*. In other words, a semantics fails to be character compositional if it assigns different characters to two complex expressions that have the same syntactic structures and pairwise synonymous constituents<sup>28</sup>. In other words, a semantics fails to be character compositional if substitution

---

27 A word on notation: from now on I'll distinguish between semantics that assigns linguistic meaning to expressions and semantics that assigns content (intensions) to expression-context pairs. I'll use  $\mathbf{I}^*$  for the first type of semantics and  $\mathbf{I}$  for the second type of semantics.

28 Expressions with identical characters are synonyms. In what follows I will use the following definition of synonymy (both for simple and complex expressions): any two expressions are synonymous iff they have identical characters. In order to ease the exposition I borrow from Szabó the term "pairwise synonymous constituents". Two complex expressions  $S_1$  and  $S_2$  have pairwise synonymous constituents iff (i) they have the same syntactic structure and (ii) for any constituent  $e$  of  $S_1$  its syntactic position in  $S_2$  is occupied by a synonymous expression. For example, (1) and (2) are expressions with pairwise synonymous constituents since "attorneys" and "lawyers" are synonymous and they occupy the same syntactic position in (1) and (2). But (3) "John admires an attorney" and (4) "A lawyer admires John" are not expressions with pairwise synonymous constituents even though "lawyer" and "attorney" are synonymous. Since the notion of expressions with pairwise synonymous constituents is similar with Carnap's notion of intensional isomorphism (Carnap, 1947, 56-57) I should eliminate a potential source of misunderstanding. I'm interested in this notion for expository purposes only. Carnap, on the other hand, was interested in certain applications of it. He takes the notion of intensional isomorphism to be a definition of synonymy for complex expressions, while I don't – I stick with identity of character as definition of synonymy. As far as I'm concerned any two complex expressions that have pairwise synonymous constituents are themselves synonymous. But there are complex expressions that are synonymous and don't have pairwise synonymous constituents because they have different syntactic structures (e.g. (4) "A lawyer admires John" is synonymous with (4') "John is admired by a lawyer"). On the other hand according to Carnap two complex expressions are synonymous if and only if they have pairwise synonymous constituents (i.e. are intensionally isomorphic). Thus (4) and (4') are not synonymous according to his account.

of synonyms is not meaning preserving in that semantics. Obviously not only semantics' that assigns characters can satisfy or fail character compositionality, but their associated un-schönfinkelized semantics' (that assigns intensions to expression-context pairs) can also satisfy or fail to satisfy character compositionality. For example a semantics that assigns contents to expression-context pairs fails to be character compositional if there are some expression  $e_i, e_j, e_n$ , operation  $\alpha$  and contexts  $C^*$  such that (a) for any context  $C$ ,  $\mathbf{I}(e_i, C) = \mathbf{I}(e_n, C)$  and (b) there is a context  $C^*$  such that  $\mathbf{I}(\alpha(e_i, e_j), C^*) \neq \mathbf{I}(\alpha(e_i, e_n), C^*)$ . According to (a)  $e_j$  and  $e_n$  have identical character, while according to (b)  $\alpha(e_i, e_j)$  and  $\alpha(e_i, e_n)$  have different characters. Any semantics that matches the failure condition for compositionality of character fails to be character compositional.

For example a semantics of English fails to be character compositional if it assigns the same character to “attorneys” and to “lawyers”, but assigns different characters to

(1) “Attorneys are rich”

and

(2) “Lawyers are rich”.

Just as there are semantics' that can fail character compositionality so there are semantics' that trivially satisfy character compositionality. A semantics trivially satisfies character compositionality if for any expressions  $e_i, e_j, e_n$ , and any syntactic operation  $\alpha$  either (a) or (b) is false. For example, a semantics trivially satisfies character compositionality if for any expressions  $e_j, e_n$ ,  $\mathbf{I}^*(e_j) \neq \mathbf{I}^*(e_n)$ . This is a semantics that never allows for synonyms: there are no two expressions  $e_j, e_n$  such that  $\mathbf{I}^*(e_j) = \mathbf{I}^*(e_n)$ . Also a semantics trivially satisfies character compositionality if for any operation  $\alpha$  and any arbitrarily complex expressions  $\alpha(e_i, e_j)$  and  $\alpha(e_i, e_n)$ ,  $\mathbf{I}^*(\alpha(e_i, e_j)) = \mathbf{I}^*(\alpha(e_i, e_n))$ . This is a semantics that assigns the same character to all complex expressions that have the same syntactic structure. Why is this important? For two reasons. One has to do with the expected explanatory benefits of compositionality. If a semantics satisfies compositionality in a trivial manner, then compositionality is explanatory idle with respect to any properties of that semantics. It should be obvious that no semantics that trivially satisfies compositionality can be a descriptively adequate semantics of natural languages.

Notice that if the principle of character compositionality holds then if replacing one expression  $e_i$  with its synonym  $e_j$  results in a meaningful complex expression, then the two complex expressions must have the same character. In other words, this principle of compositionality entails that if two meaningful complex expressions  $S_1$  and  $S_2$  are such that  $S_2$  is the result of replacing a constituent of  $S_1$  with a synonym, then  $S_1$  and  $S_2$  are also synonyms.

Under the assumption that each constituent of a meaningful complex expression is itself meaningful we can formulate the principle of character compositionality in an alternative manner: a semantics is character compositional iff for any simple expressions  $e_i$ ,  $e_j$ ,  $e_n$  and syntactic operation  $\alpha$  if  $\mathbf{I}^*(e_j) = \mathbf{I}^*(e_n)$  and if  $\alpha(e_i, e_j)$  and  $\alpha(e_i, e_n)$  are both meaningful then  $\mathbf{I}^*(\alpha(e_i, e_j)) = \mathbf{I}^*(\alpha(e_i, e_n))$ . In other words, a semantics is character compositional iff given any two meaningful complex expressions  $S_1$  and  $S_2$  where  $S_2$  is obtained from  $S_1$  by replacing a constituent of  $S_1$  with a synonymous expression, the character of  $S_1$  is identical with the character of  $S_2$ <sup>29</sup>. Many authors formulate the principle of compositionality in terms of substitution of meaning-equivalents (e.g. Carnap 1947, 51-52), and many of them talk about there being a function version and a substitution version of compositionality. Strictly speaking the two formulations are equivalent only under the assumption that each constituent of a meaningful complex expression is itself meaningful.

This principle can be employed both by semantics of languages with no context sensitivity, (that is languages in which the character of each expression is a constant function) and by semantics of languages with context-sensitive expressions (one in which the characters of some of its expressions are a non-constant function).

### 3.2 Compositionality for Content

For a semantics that pairs expressions with their meanings the formulation of the principle is very much straightforward. But for a semantics that assigns meanings to

---

<sup>29</sup> Another way to put it is this: a semantics is character compositional only if any meaningfulness-preserving substitution of an expression with its synonym is also a meaning-preserving substitution.

expression-context pairs, the formulation is less so, since we need to take into account the role of the context in the interpretation of complex expressions. So the principle of compositionality needs to be slightly modified so to make room for the contribution of the context to the interpretation of expressions. Since there are two broad positions of how context contributes to the interpretation of complex expressions, there are two different non-equivalent ways to spell out the principle of compositionality for a semantics that assigns meanings to expression-context pairs.

We can modify the principle of compositionality in two different ways so to take into account the role of the context. One is stronger than the other, and given that the terminology is also entrenched in the literature, I will call one, *strong compositionality* and the other *weak compositionality*<sup>30</sup>.

### 3.2.1. Strong Compositionality for Content

According to strong compositionality the content of a complex expression relative to a context C is determined by the contents of its constituents relative to context C and the way they are syntactically combined.

A more precise formulation of strong compositionality is the following: a semantics **I** is strongly compositional iff for every syntactic rule  $\alpha$  there is a function  $f$  such that for any two expressions  $e_i, e_j$  and for any context C if  $\alpha(e_i, e_j)$  is meaningful at C :

$$\mathbf{I}(\alpha(e_i, e_j), C) = f(\alpha, \mathbf{I}(e_i, C), \mathbf{I}(e_j, C))$$

A semantics fails to be strongly compositional if for some expressions  $e_i, e_j$ , syntactic operation  $\alpha$  and some contexts  $C_1, C_2$  (c) and (d) hold:

$$(c) \mathbf{I}(e_i, C_1) = \mathbf{I}(e_i, C_2) \text{ and } \mathbf{I}(e_j, C_1) = \mathbf{I}(e_j, C_2)$$

and

---

<sup>30</sup> A historical: traditionally when theorists gave a formulation of compositionality for semantics that assigns meaning to expressions in context they have done it in the form of strong compositionality (see for example Kaplan 1989: 507). Weak compositionality is a later addition in the semantic landscape. Weak compositionality was motivated by alleged failures to give a strongly compositional treatment to certain expressions. In the third and fourth chapter I will discuss at length the motivation behind the idea of weak compositionality.



(d)  $\mathbf{I}(\alpha(e_i, e_j), C_1) \neq \mathbf{I}(\alpha(e_i, e_j), C_2)$ .

I will call the above statement *the failure condition for strong compositionality*. In the semantic literature there are several proposed semantics for English (or fragments of English) that fail to be strongly compositional. In the next section I will discuss one such semantics for the purposes of illustration.

On the other hand a semantics trivially satisfies strong compositionality if for any expressions  $e_i, e_j$ , any syntactic operation  $\alpha$  and any contexts  $C_1, C_2$  either (c) or (d) are false. For example, a semantics trivially satisfies strong compositionality if for any expression  $e_i$  and any contexts of utterance  $C_1$  and  $C_2$ ,  $\mathbf{I}(e_i, C_1) = \mathbf{I}(e_i, C_2)$ . This is a semantics that assigns values that vary with each context of utterance. In other words, it is a semantics that for any expression  $e_i$ , there are no contexts  $C_1$  and  $C_2$  such that the content of  $e_i$  at the two contexts coincides. A semantics can also trivially satisfy strong compositionality if for any complex expression  $S$ , and any contexts  $C_1$  and  $C_2$ ,  $\mathbf{I}(S, C_1) = \mathbf{I}(S, C_2)$ . This is a semantics in which all complex expressions are context insensitive (i.e. their content is independent of the context). To put it intuitively a semantics that assigns to each complex expression a context-independent semantic content trivially satisfies strong compositionality.

A consequence of strong compositionality, relevant for our later discussion, is that the content of a complex expression “depends on the context only in so far as the [contents] of its constituents do” (Szabó 2001, 122). If the content of complex expression is context dependent this should be traceable to at least one of its simple (atomic) constituents.

### 3.2.2. Weak Compositionality for Content

According to weak compositionality the content of a complex expression relative to a context  $C$  is a function of the contents that its constituents have at  $C$  *and of  $C$  itself*.

A more precise formulation of weak compositionality is the following: a semantics  $\mathbf{I}$  is

weakly compositional iff: for every syntactic rule  $\alpha$  there is a function  $f$  such that for any expressions  $e_i, e_j$  and for any context  $C$  if  $\alpha(e_i, e_j)$  is meaningful then

$$\mathbf{I}(\alpha(e_i, e_j), C) = f(\alpha, \mathbf{I}(e_i, C), \mathbf{I}(e_j, C), C)$$

A semantics fails to be weakly compositional if for some expressions  $e_i, e_j, e_n$ , syntactic operation  $\alpha$  and context  $C$ :

$$(e) \mathbf{I}(e_j, C) = \mathbf{I}(e_n, C)$$

and

$$(f) \mathbf{I}(\alpha(e_i, e_j), C) \neq \mathbf{I}(\alpha(e_i, e_n), C)$$

I will call the above statement *the failure condition for weak compositionality*. That is, if relative to a context  $C$  any two expressions  $e_j$  and  $e_n$  are assigned the same content  $\mathbf{I}(e_j, C) = \mathbf{I}(e_n, C)$ , but two complex expressions  $\alpha(e_i, e_j)$  and  $\alpha(e_i, e_n)$ , where the second expression is the result of replacing  $e_j$  with  $e_n$  within the first expression, are assigned different contents relative to  $C$ , this brings about a failure of weak compositionality. I will give an example of a semantics that fails to be weakly compositional in the next section.

A semantics trivially satisfies weak compositionality if for any expressions  $e_i, e_j$ , syntactic operation  $\alpha$  and context  $C$  either (e) or (f) is false. That is, a semantics is trivially weakly compositional if for any two expressions  $e_i$  and  $e_j$  and any context  $C$ ,  $\mathbf{I}(e_i, C) \neq \mathbf{I}(e_j, C)$ . In other words there are no two expressions and context  $C$  such that the two expressions agree on content at  $C$  (i.e. they have the same content relative to  $C$ ). A semantics can also trivially satisfy weak compositionality if for any syntactic rule  $\alpha$  and any complex expressions  $S_1$  and  $S_2$  built by  $\alpha$ , and any context  $C$ ,  $\mathbf{I}(S_1, C) = \mathbf{I}(S_2, C)$ . This is a semantics that assigns to any two complex expressions with identical syntactic structure equivalent contents relative to any context of utterance<sup>31</sup>.

Notice that according to weak compositionality (as opposed to strong compositionality), context is taken as an extra argument of the semantic operation and *the contribution that this extra argument makes is not vacuous* (Pagin: 2005, 55). In fact a semantics that satisfies weak compositionality and in which the extra contribution of the

---

31 That is to say that according to such a semantics any two expressions with the same syntactic structure are synonymous.

context is *always* vacuous, also satisfies strong compositionality. Moreover for such a semantics it holds true what Szabó calls the context principle, namely that the content of an expression depends on the context only insofar as the contents of its constituents do. Theorists are interested in weak compositionality precisely because it allows context to have a *non-vacuous* contribution to the interpretation of the complexes that is not reducible to its contribution to the interpretation of the constituents and because it allows that different contexts contribute in different ways to the interpretation of complex expressions over and above their contribution to the interpretation of the parts. Than a weakly compositional semantics in which the extra contribution of the context is constant (i.e. every context makes the same contribution to the interpretation of complexes) is also of no interest to theorists that advocate weak compositionality. A semantics in which the extra contribution of the context is constant is one which also satisfies strong compositionality and one in which the context sensitivity of complex expressions is traceable to the context-sensitivity of the parts (although it need not be that every material in the content of complex expressions is traceable to the content of the parts or the semantic effect of syntax). For example a language (call it English\*) with the same vocabulary and syntax as English only that it differs from English in that for any sentence S\* (of English\*) its content at any context C is the conjunction of the content of its English counterpart S at C and the proposition that  $1+1=2$ . If English is strongly compositional then so is English\*. In English\* not every material in the content of the complexes is traceable to the content of its constituents (e.g. the second conjunct of the content,  $1+1=2$ , is not), but any contextual variation in the content of complexes *is* traceable to the content of its constituents.

A consequence of weak compositionality is that context sensitivity of complex expressions not necessarily traceable to the context-sensitivity of some of constituents. As some advocates of weak compositionality (e.g. Pelletier (2003), Recanati (2010)) emphasize weak compositionality allows that context determines the content of a complex expression over and above determining the contents of its constituents.

## 4. Relations Between Principles of Compositionality<sup>32</sup>

The following relations hold between the principles of compositionality introduced above. If a semantics satisfies strong compositionality then it also satisfies weak and character compositionality. A semantics can satisfy weak compositionality but fail to satisfy strong compositionality. If it satisfies weak compositionality then it satisfies character compositionality. A semantics can satisfy character compositionality but fail to satisfy weak compositionality. If a semantics fails to satisfy strong compositionality it then it will fail to satisfy weak compositionality and character compositionality.

We can represent the relation between them in the following way, where the left to right arrow represents entailment:

Strong compositionality  $\rightarrow$  Weak compositionality  $\rightarrow$  Character compositionality.

### 4.1 Strong and Weak Compositionality

That strong compositionality entails weak compositionality should be obvious. Any semantics that satisfies strong compositionality is (or can be represented as) a semantics that satisfies weak compositionality in which the extra-contextual contribution is either always vacuous or always constant.

On the other hand, given that according to weak compositionality, the composition function takes the context as an extra argument and that the contribution of the extra argument *can be* non-vacuous, then it is obvious that there are possible semantics that are weakly compositional but fail to be strongly compositional. A semantics that assigns to complex expressions different contents at different contents but assigns constant contents to their constituents fails to be strongly compositional. Nevertheless such a semantics *can be* weakly compositional. I will illustrate this point with the help of a proposed semantics (by Pelletier 2003) for quantified noun phrases, that fails to satisfy strong compositionality but satisfies weak compositionality.

---

<sup>32</sup> Westerståhl (2012) reports the same results but the proofs differ.

Intuitively different utterances of (3) are true even though not every student in the universe failed the exam.

(3)“Every student failed the exam”

Moreover different utterances of (3) can have different truth conditions. For example, if uttered by a logic professor (while talking about her teaching) we will judge (3) to be true just if every student in her logic class failed the exam. But if (3) is uttered by a math professor we will judge it to be true just if every student in her math class failed the exam.

One way to capture this intuition is to claim that different utterances of (3) express different contents; that the extension of the quantified noun phrase “every student” is restricted to a contextually specified group *and that the context of utterance is what restricts the quantifier domain*. For example if a logic professor utters (3) (while talking about her teaching) we might take (3) to be true iff every student in her logic class failed the exam. While if a math professor utters (3) we might take it to be true iff every student in her math class failed the exam.

There are several ways to model how context restricts the domain of the quantifier phrase<sup>33</sup>. Most accounts proceed in the following way: relative to different contexts  $C_1$  and  $C_2$  (3) expresses different contents: at  $C_1$  (3) expresses the content that every student in the logic class failed the exam, while at  $C_2$  it expresses the content that every student in the math class failed the exam. We can model the content (3) relative to  $C_1$  as a function from worlds to truth values, such that it yields true if every student in the logic class failed the exam at world  $w$  and yields false otherwise, and its content at  $C_2$  as a function that yields true if every student in the math class failed the exam at world  $w$ , and yields false otherwise. Then the domain of quantification is part of the content of (3) and the provided by the context. A second step is to claim that the restriction of the domain of quantification is part of the content of the quantified noun phrase “every student”.

A semantics that fails to be strongly compositional but can be weakly compositional is one that assigns to “every student” contents that vary across contexts but assigns stable

---

<sup>33</sup> For a thorough overview of the different options to account for quantifier domain restriction see Stanley and Szabó 2000.

contents to its constituents “every” and “student”. More exactly, a semantics fails to be strongly compositional but it can be weakly compositional if there is domain restriction on the entire quantified noun phrase “every student”, but not on any of its constituents. and this is spelled out in terms of having the domain part of the content of the quantified noun phrase.

Here are the kind of assignments that such a semantics must make. I will not go into the details of the theory (for example, how the contextual material is represented). What is important here is where in the analysis of (3) the domain restriction comes in. I will help myself of the lambda notation to keep things clear.

(i)  $\mathbf{I}(\textit{Every student failed the exam}, C_1) = \lambda w. [|\textit{Every student failed the exam}|]^{w,C_1} =$  a function  $f$  such that for any possible world  $w$ ,  $f(w)$  returns truth iff every individual that is a student in the logic class failed the exam in world  $w$ .

(ii)  $\mathbf{I}(\textit{Every student failed the exam}, C_2) = \lambda w. [|\textit{Every student failed the exam}|]^{w,C_2} =$  a function  $f$  such that for any possible world  $w$ ,  $f(w)$  returns truth iff every individual that is a student in the math class failed the exam in world  $w$ .

Again we take contexts  $C_1$  and  $C_2$  as quadruples consisting of an agent, world, time and location such that  $C_1$  and  $C_2$  differ *only* in their respective agent.

(iii)  $\mathbf{I}(\textit{every student}, C_1) = \lambda w. [|\textit{every student}|]^{w,C_1}$ , where  $[|\textit{every student}|]^{w,C_1}$  is a function from a first-order property to a truth-value that yields truth iff every student in the logic class satisfies that first-order property. In plain words, “every student” states in this context is that the predicate of the sentence is true of all individuals who are students in the logic class. So its content can be modeled as a function from possible worlds to a second-order property that yields true iff the predicate is true of every individual that is a student in the logic class.

(iv)  $\mathbf{I}(\textit{every student}, C_2) = \lambda w. [|\textit{every student}|]^{w,C_2}$ , where  $[|\textit{every student}|]^{w,C_2}$  is a function from a first-order property to a truth-value that yields truth iff every student in the math class satisfies that first-order property.

(v)  $\mathbf{I}(\textit{student}, C_1) = \lambda w. [|\textit{student}|]^{w,C_1}$ , where  $[|\textit{student}|]^{w,C_1}$  = the set of individuals that are

students at  $w$

(vi)  $\mathbf{I}(\textit{student}, C_2) = \lambda w. [\textit{student}]^{w, C_2}$ , where  $[\textit{student}]^{w, C_2}$  = the set of individuals that are students at  $w$

(vii)  $\mathbf{I}(\textit{every}, C_1) = \lambda w. [\textit{every}]^{w, C_1}$  where  $[\textit{every}]^{w, C_1}$  = is a function  $f$  from first-order properties to second-order properties such that for any property  $p$ ,  $f(p)$  = the second order property that has  $p$  as a subset.

(viii)  $\mathbf{I}(\textit{every}, C_2) = \lambda w. [\textit{every}]^{w, C_2}$  where  $[\textit{every}]^{w, C_2}$  = the same as above (a function  $f$  from first-order properties to second-order properties such that for any property  $p$ ,  $f(p)$  = the second order property that has  $p$  as a subset.)

Obviously such a semantics fails to be strongly compositional since its assignments match the failure condition for strong compositionality: for some expressions  $e_i$ ,  $e_j$ , operation  $\alpha$  and contexts  $C_1$  and  $C_2$  such that  $\mathbf{I}(e_i, C_1) = \mathbf{I}(e_i, C_2)$  and  $\mathbf{I}(e_j, C_1) = \mathbf{I}(e_j, C_2)$  and  $\mathbf{I}(\alpha(e_i, e_j), C_1) \neq \mathbf{I}(\alpha(e_i, e_j), C_2)$ .

According to (iii) and (iv), the content of the quantified noun phrase “every student” varies across the two contexts of utterance, while according to (v)-(viii) both its constituents have unvarying contents at the two contexts. The common noun “student” gets assigned the same content at  $C_1$  and  $C_2$ , namely the first-order property of being a student, and likewise for the determiner “every”: its content at  $C_1$  and  $C_2$  coincide. Put somewhat differently, the quantified noun phrase contains a contextually restricted domain of quantification (the students in the logic class at  $C_1$ , and the students in the math class at  $C_2$ ) although none of its constituents contain the restriction.

Nevertheless such a semantics *can* be weakly compositional since these assignments do not match the failure condition for weak compositionality. In other words, given that weak compositionality lets the context be taken as an extra argument by the composition function Then if its contribution is non-vacuous complex expressions can have different contents at different contexts even if none of their constituents vary their content across those contexts of utterance. The content of a complex expression relative to a context is a function of the contents of its constituents relative to that context, its syntactic structure *and of the context itself*.

## 4.2 Weak Compositionality and Character Compositionality

There are semantics that are character compositional but not weakly compositional.

For reasons of clarity I'll start by rehearsing how weak compositionality and character compositionality fail.

A semantics fails to be weakly compositional if the following obtains: for some expressions  $e_i, e_j, e_n$ , syntactic rule  $\alpha$  and context of utterance  $C$ :

$$(e) \mathbf{I}(e_j, C) = \mathbf{I}(e_n, C) \text{ and } (f) \mathbf{I}(\alpha(e_i, e_j), C) \neq \mathbf{I}(\alpha(e_i, e_n), C)$$

That is, if relative to a context  $C$  two expressions  $e_j$  and  $e_n$  are assigned the same content  $\mathbf{I}(e_j, C) = \mathbf{I}(e_n, C)$ , but two complex expressions  $\alpha(e_i, e_j)$  and  $\alpha(e_i, e_n)$  are assigned different contents relative to  $C$ , then such a semantics is not weakly compositional.

By schönfinkelizing  $\mathbf{I}$ ,  $\mathbf{I}(e_j, C) = \mathbf{I}^*(e_j)(C)$ . So we can re-write the failure condition in the following way: for some expressions  $e_i, e_j, e_n$ , syntactic rule  $\alpha$  and context of utterance  $C$ :

$$(e) \mathbf{I}^*(e_j)(C) = \mathbf{I}^*(e_n)(C) \text{ and } (f) \mathbf{I}^*(\alpha(e_i, e_j))(C) \neq \mathbf{I}^*(\alpha(e_i, e_n))(C)$$

A semantics fails to be character compositional if the following obtains: for some expressions,  $e_i, e_j, e_n$ , and syntactic rule  $\alpha$ :

$$(a) \mathbf{I}^*(e_j) = \mathbf{I}^*(e_n) \text{ and } (b) \mathbf{I}^*(\alpha(e_i, e_j)) \neq \mathbf{I}^*(\alpha(e_i, e_n))$$

Let's consider a language  $L$  (a fragment of English) for which a weakly compositional semantics  $\mathbf{I}$  can be given. Assume that that for the two sentences  $(\Phi)$  "I am hungry" and  $(\Psi)$  "Adrian is hungry" of  $L$  and context  $C_1$  (where  $C_1$  is such that Adrian is the agent of  $C_1$ )  $\mathbf{I}$  makes the following assignments.

If  $\Phi$  uttered at  $C_1$  it says about agent of  $C_1$  and time of  $C_1$  that the former is hungry at the later. That is  $\Phi$  is true at  $C_1$  iff the agent of  $C_1$  is hungry at the time of  $C_1$ . Since Adrian is the agent of  $C_1$ ,  $\Phi$  is true at  $C_1$  iff Adrian is hungry at  $C_1$ . The content of  $\Phi$  at  $C_1$  is or can be modeled as a function from world-time pairs to truth values that it yields true at for those  $\langle w, t \rangle$  such that Adrian is hungry in  $w$  at  $t$ , and yields false otherwise. Sentence  $\Psi$  at  $C_1$  gets assigned a content that is a function from world-time pairs to truth values that it yields true at for those  $\langle w, t \rangle$  such that Adrian is hungry in  $w$  at  $t$ , and yields false otherwise. Relative



to  $C_1$ ,  $\Psi$  and  $\Phi$  express the same content:  $\mathbf{I}(\Phi, C_1) = \mathbf{I}(\Psi, C_1)$ .

Since by assumption such a semantics is weakly compositional the contents of  $\Phi$  and  $\Psi$  relative to  $C_1$  is a function of the contents of their respective constituents at  $C_1$ , their syntactic structure and  $C_1$  itself.

Now let's consider an extension  $M$  of the initial language  $L$ .  $M$  contains all expressions of  $L$  plus the sentential operator "At context  $C_2$ " ( $\text{At}C_2$  for short) and sentences constructed with the help of this operator. Thus, for any sentence  $S$  if  $S$  is a sentence of  $L$  then  $S$  is also a sentence of  $M$ . And for any sentence  $S$  of  $L$  there is a sentence of the form  $\alpha(\text{At}C_2, S)$  in the extended language  $M$ <sup>34</sup>. We give a semantics  $\mathbf{I}_M$  to  $M$  that preserves or inherits the assignments of semantics  $\mathbf{I}$  of  $L$ . For any sentence  $S$  of  $L$  the semantic function  $\mathbf{I}_M$  of  $M$  agrees with the semantic function  $\mathbf{I}$  of  $L$ :  $\mathbf{I}(S) = \mathbf{I}_M(S)$ . In other words, the semantics  $\mathbf{I}_M$  of  $M$  is the same as semantics  $\mathbf{I}$  of  $L$  for all sentences that do not contain the operator  $\text{At}C_2$ .

The operator  $\text{At}C_2$  is a function such that for any sentence  $S$  and any context  $C$  it takes the content expressed by  $S$  at  $C$  (i.e.  $\mathbf{I}_M(S, C)$ ) and yields the content expressed by  $S$  at  $C_2$  (i.e.  $\mathbf{I}_M(S, C_2)$ )<sup>35</sup>. In other words, for any context  $C$  a sentence of the form  $\alpha(\text{At}C_2, S)$  is true relative to  $C$  iff  $S$  is true at  $C_2$ . Which means that for any context  $C$ ,  $\mathbf{I}_M(\alpha(\text{At}C_2, S), C) = \mathbf{I}_M(S, C_2)$

It is easy to show that semantics  $\mathbf{I}_M$  of the extended language  $M$  is not weakly compositional. Take the two sentences  $(\Phi)$  "I am hungry" and  $(\Psi)$  "Adrian is hungry" and build two sentences with the help of the sentential operator  $\text{At}C_2$ . We get the following sentences in the extended language:

(A) "At context  $C_2$  I am hungry"

---

34 Since  $\text{At}C_2$  is a sentential operator I could write the sentence obtained by applying the operator to a sentence  $S$  more elegantly as  $\text{At}C_2(S)$ . But to keep in line with the formalism introduced and to keep things clear with respect to compositionality I'll use the slightly more cumbersome notation and write the sentence thus obtained as  $\alpha(\text{At}C_2, S)$ .

35 Obviously the  $\text{At}C_2$  operator is what Kaplan calls a monster (see Kaplan 1989, 510-511). Whether English contains expressions that are monsters and whether such expressions can be added to English is still a matter of debate. Kaplan claims that English neither contains such expressions nor they can be added to it. Israel and Perry (1996) and Schelnker (2003) argue otherwise. For this reason I use the technical expression  $\text{At}C_2$  and not any of the expressions that some of the latter semanticists think are monsters. Obviously the claims about the truth values of sentences containing  $\text{At}C_2$  rely on stipulation and not on intuitions that speakers might have about them.

(B) “At context  $C_2$  Adrian is hungry”

Suppose that Adrian is the agent of  $C_1$  and suppose that the agent of  $C_2$  is not Adrian.

(A) is true at  $C_1$  iff the speaker of  $C_2$  is hungry at the world and time of  $C_2$ . And (B) is true at  $C_1$  iff Adrian is hungry at the world and time of  $C_2$ .

The content expressed by (A) at  $C_1$  is different than the content expressed by (B) at  $C_1$ :  
 $\mathbf{I}_M(\alpha(\text{At}C_2, \Phi), C_1) \neq \mathbf{I}_M(\alpha(\text{At}C_2, \Psi), C_1)$

Given that  $M$  is an extension of the initial language  $L$ , and its semantics  $\mathbf{I}_M$  preserves the assignments of the semantics  $\mathbf{I}$  for expressions of  $L$  since  $\mathbf{I}(\Phi, C_1) = \mathbf{I}(\Psi, C_1)$  then  $\mathbf{I}_M(\Phi, C_1) = \mathbf{I}_M(\Psi, C_1)$ .

$\mathbf{I}_M$  is not weakly compositional since for some expressions  $\Phi$  and  $\Psi$ , syntactic rule  $\alpha$  and context  $C_1$ ,  $\mathbf{I}_M(\alpha(\text{At}C_2, \Phi), C_1) \neq \mathbf{I}_M(\alpha(\text{At}C_2, \Psi), C_1)$  and  $\mathbf{I}_M(\Phi, C_1) = \mathbf{I}_M(\Psi, C_1)$ , which match the failure condition for weak compositionality. ■

The extended language can, nevertheless, receive a character compositional semantics. Or, in other words we can define a function  $\mathbf{I}^*_M$  associated with the  $\mathbf{I}_M$  semantics such that  $\mathbf{I}^*_M$  is character compositional: for any sentence  $s$  and any context  $C$ ,  $\mathbf{I}_M(S, C) = \mathbf{I}^*_M(S)(C)$  (again  $\mathbf{I}_M$  assigns functions from world-time pairs to extensions (i.e. contents) and  $\mathbf{I}^*_M$  assigns functions from contexts to contents).

We need to show that the following holds for  $\mathbf{I}^*_M$ : for any syntactic rule  $\alpha$ , there is a function  $f$  such that any arbitrary expressions  $e_i, e_j$  of  $M$  if  $\mathbf{I}^*_M(\alpha(e_i, e_j))$  is defined then  $\mathbf{I}^*_M(\alpha(e_i, e_j)) = f(\alpha, \mathbf{I}^*_M(e_i), \mathbf{I}^*_M(e_j))$ . Or, in other words, we need to show that there are no expressions  $e_i, e_j, e_n$ , and syntactic rule  $\alpha$  such that  $\mathbf{I}^*_M(e_j) = \mathbf{I}^*_M(e_n)$  and  $\mathbf{I}^*_M(\alpha(e_i, e_j)) \neq \mathbf{I}^*_M(\alpha(e_i, e_n))$ .

The  $\mathbf{I}_M$  semantics for  $M$  holds the following as true: for any context  $C$ ,  $\mathbf{I}_M(\alpha(\text{At}C_2, \Phi))(C) = \mathbf{I}_M(\Phi, C_2)$  (by definition of the  $\text{At}C_2$  operator).

We can write this in the  $\mathbf{I}^*_M$  semantics as: for any  $C$ ,  $\mathbf{I}^*_M(\alpha(\text{At}C_2, \Phi))(C) = \mathbf{I}^*_M(\Phi)(C_2)$ . Thus for (A) of  $M$ ,  $\mathbf{I}^*_M$  makes the following assignment:

(1) For any  $C$ ,  $\mathbf{I}^*_M(\alpha(\text{At}C_2, \Phi))(C) = \mathbf{I}^*_M(\Phi)(C_2)$

and (B) has the following  $\mathbf{I}^*_M$  assignment

(2) For any  $C$ ,  $\mathbf{I}^*_M(\alpha(\text{At}C_2, \Psi))(C) = \mathbf{I}^*_M(\Psi)(C_2)$

We can show that  $\mathbf{I}^*_M$  character compositional by *reductio*. Let's assume that  $\mathbf{I}^*_M$  for  $M$  is not character compositional. Then for some expressions  $\Phi$  and  $\Psi$   $\mathbf{I}^*_M$  makes the following assignments which match the failure condition for character compositionality:

(3)  $\mathbf{I}^*_M(\alpha(\text{At}C_2, \Phi)) \neq \mathbf{I}^*_M(\alpha(\text{At}C_2, \Psi))$

(4)  $\mathbf{I}^*_M(\Phi) = \mathbf{I}^*_M(\Psi)$

From (4) we get:

(5) For any  $C$ ,  $\mathbf{I}^*_M(\Phi)(C) = \mathbf{I}^*_M(\Psi)(C)$

From (5) we get:

(6)  $\mathbf{I}^*_M(\Phi)(C_2) = \mathbf{I}^*_M(\Psi)(C_2)$

From (1), (2), and (6) we get:

(7) For any  $C$ ,  $\mathbf{I}^*_M(\alpha(\text{At}C_2, \Phi))(C) = \mathbf{I}^*_M(\alpha(\text{At}C_2, \Psi))(C)$

From (3) and the definition of  $\mathbf{I}^*_M$  we get:

(8) There is a  $C$ , such that  $\mathbf{I}^*_M(\alpha(\text{At}C_2, \Phi))(C) \neq \mathbf{I}^*_M(\alpha(\text{At}C_2, \Psi))(C)$

But (8) contradicts (7), therefore  $\mathbf{I}^*_M$  is character compositional - i.e.  $\neg((3)\&(4))$ . ■

So, to conclude, failure of weak compositionality does not entail failure of character compositionality: there are semantics that are character compositional without being weakly compositional.

It can be shown that if a semantics is weakly compositional then it is also character compositional (i.e. weak compositionality entails character compositionality).

This is a *proof by contraposition*. It is easy to prove that if a semantics fails to be character compositional then it also fails to be weakly compositional (i.e. failure of character compositionality entails failure of weak compositionality) which is *the contrapositive* of the claim that if a semantics satisfies weak compositionality it also satisfies character compositionality.

A semantics fails to be character compositional if for some expressions  $e_i$ ,  $e_j$ , and  $e_n$  and syntactic rule  $\alpha$  the following assignments hold:

(a)  $\mathbf{I}^*(e_j) = \mathbf{I}^*(e_n)$  and (b)  $\mathbf{I}^*(\alpha(e_i, e_j)) \neq \mathbf{I}^*(\alpha(e_i, e_n))$ .

Notice that:

(a) entails (a'):  $\forall C, \mathbf{I}^*(e_j)(C) = \mathbf{I}^*(e_n)(C)$  and

(b) entails (b'):  $\exists C^*$  such that  $\mathbf{I}^*(\alpha(e_i, e_j))(C^*) \neq \mathbf{I}^*(\alpha(e_i, e_n))(C^*)$ .

The conjunction of (a') and (b') entails failure of weak compositionality: there are some expressions  $e_i, e_j, e_n$ , syntactic rule  $\alpha$  and a context  $C^*$  such that  $\mathbf{I}^*(e_j)(C^*) = \mathbf{I}^*(e_n)(C^*)$  and  $\mathbf{I}^*(\alpha(e_i, e_j))(C^*) \neq \mathbf{I}^*(\alpha(e_i, e_n))(C^*)$ .■

In conclusion, if a semantics satisfies weak compositionality for content then it also satisfies character compositionality.

#### 4. Summing up

In this chapter I did several things. First, I gave a mathematically precise formulation of the principle of compositionality. Then, I discussed how the principle can accommodate context-sensitivity. There are two popular strategies for modeling the context-sensitivity of natural languages. One is to assign semantic values to sentences relative to contexts, where the semantic values assigned are entities which are amenable to evaluation in terms of truth and falsity. The other strategy is to assign semantic values directly to expressions, and the semantic values assigned are functions from contexts to entities which, in their turn, are evaluable in terms of truth and falsity. If the semantics is defined over expression-context pairs then we need to modify the formulation of the principle of compositionality so to take the contribution of the context into account. I also discussed three principle of compositionality, the constraints they place on semantic theories and the relations among them. These three principles can be ordered by strength: there is a strong compositionality principle for content, a weak compositionality principle for content, and finally, the weakest, character compositionality.



## **CHAPTER 3: Context Sensitivity and Compositionality.**

### **1. Introduction**

The picture about compositional semantic theories that we get from chapter one is the following. Words (that is, simple expressions) have meanings and their meanings are the linguistic conventions associated with them. These linguistic conventions are or determine a rule of correct use for the associated expression. Words combine in certain ways to form larger expressions and the way they combine can be specified in the syntactic analysis of the language. Some of these larger expressions formed by combining words according to the syntactic rules are meaningful, and their meaning is a function of the meaning of their constituents and syntactic structure. When it comes to sentences, their linguistic meaning determines a set of conditions under which their literal use is true – a set of conditions such that a literal utterance of that sentence is true iff those conditions are satisfied. For some sentences of natural languages the conditions under which their utterances are true vary with the context of utterance, thus their literal meaning determines their truth conditions only relative to a context. This is, in a nutshell, the standard (or traditional) view of formal semantic theories: that one can give a theory that generates truth-value predictions by simply “operating a calculus according to definite rules” as Wittgenstein put it (1953: §81).

But right from their beginnings formal semantic theories were met with skepticism. The skepticism concerns mainly the claim that expressions have, or at least that a semanticist can specify, an associated rule that determines all their correct literal uses. According to skeptics, speakers use expressions to refer to things and properties in various ways, reflecting their immediate interests and situations. It is wrong to believe that the linguistic meaning of an expression will determine all of its correct uses, or, at least, that a

semanticist can specify for any expression a rule or convention that will determine its correct use for any possible situation of use. In other words, skeptics doubt the very possibility of formal semantics as described in chapter one. According to them, for virtually any declarative sentence of natural languages, its truth-conditions are determined by a myriad of unsystematic pragmatic factors.

## **2. Semantic and Pragmatic Explanations.**

Let me clarify first what goes under the labels of “semantics” and “pragmatics” and what are the claims of semantic and pragmatic theories about.

As I explained in the first chapter, formal semantic theories aim to model competence with a language by means of generating truth value predictions for sentences of a model language that stands in a certain relation (i.e the actual-language relation) with the language used by a group of speakers whose competence is modeled. If given the bridge laws described in the first chapter, the actual-language relation is satisfied, that is if the language as described by the theory *is* the actual language used by some population, then we say that the semantic theory models the linguistic competence of members of that population.

Semantics as done in this thesis, does not describes the cognitive processes causally responsible for linguistic competence. If theorists can determine empirically that a language L, as described by a semantic theory is used by a group of speakers it doesn't follow that the entities assigned by semantics or that the rules employed by semantics describe psychologically real mechanisms that are causally responsible for users' competence with L. Therefore compositionality as understood here is a property of languages as abstract objects and not a property of the psychology of language users. I take the principles of compositionality to be about meaning-relations, namely about the relation between the meaning of complex expressions and the meaning of their constituents. Principles of compositionality, as understood here, do not aim to describe some recursive mechanism realized in the psychology of language users, that is causally responsible for how speakers

understand and produce complex expressions. It might turn out that there is a convergence between formal semantic theories and psychological theories of language use, in the sense that the structures and rules posited by formal semantic theory correspond exactly to those posited by a psychological theory of language use. But even so, that doesn't make formal semantic theories, theories about psychological processes, since their statements are about truth-conditions. Statements about truth-conditions (and derivable statements about meaning-properties and meaning-relations like statements about synonymy, entailment, logical truth, rigidity, derivable from statements about truth conditions) are not about psychological processes for the very simple reason that to give truth-conditions of sentences is to specify the worldly conditions. Secondly, a theory that makes claims about meaning properties and meaning relations of a language does not determine a unique theory of the internal psychological processes and states of language users. In other words, given that semantic theories do not aim to offer a direct description of actual psychological processes that people go through in real time, they will be judged to be correct or not independently of how its rules and entities are implemented psychologically. Even if there is such a correspondence, semantic theories and psychological theories about language use, make different claims and are accountable to different types of data. A semanticist need not assume that her theory describes psychological processes and states that underlie language use. I will adopt this take on what semantics is about<sup>36</sup>.

Semantic theories are to be supplemented by pragmatic theories that seek to explain how speakers use sentences to communicate – that is, theories that seek to explain how sentences can be used in different situations to convey certain information. Pragmatic theories also play another role: they filter out those intuitions that are relevant for semantic theorizing from those that are not. Any sentence can be used in a literal way or in a non-literal way, for example in an ironic way or in a metaphoric way. An uttered sentence can be evaluated in several different ways, depending on whether the utterance was literal, or, for example, ironic. One can utter “Adrian is a fine friend” literally so to let the audience know that Adrian is a fine friend but one can also utter it ironically so to let the audience know that Adrian is *not* a fine friend. Obviously we can evaluate this in different ways,

---

36 Soames (1984, 1985) defends this way of understanding what semantics is about.



depending on whether one speaks literally or ironically. If Adrian is not a fine friend and one speaks literally then one utters a falsehood, but if one speaks ironically then one speaks the truth. To explain these different evaluations, pragmatic theories distinguish the content expressed by a sentence at a context (what is assigned by the semantics), and what is conveyed by uttering that sentence at that context. They need not be identical: what is conveyed by uttering a sentence at a context *can* be different than the semantic content of the sentence. This is the case with the ironic use of “Adrian is a fine friend”. Then the intuitions relevant for semantic theorizing, that is the intuitions that constitute the data that semantic theories aim to explain and predict, are, the intuitions about the truth-values of literal uses of sentences. Pragmatic theories must provide a series of tests that will tell those intuitions apart from other intuitions not relevant for semantic theorizing.

To sum up: neither semantics nor pragmatics, as understood here, deal with the psychological details of language use. As done in this thesis semantics is not in the business of giving a psychological model of the mechanisms by which language users pair expressions with their meanings. And pragmatics is not in the business of giving a psychological description of the processes by which language users exploit a wide range of information in order to interpret an utterance-event. Semantics is the description of a language as an abstract object whereby the expressions of the language (as individuated by syntax) are associated with aspects of the world – i.e. semantics assigns individuals and properties to simple expressions such that together with the rules of combination they determine for any sentence of the language its truth conditions. Pragmatics offers a rational reconstruction of how language users deploy sentences in a conversation so to get across information. This way of understanding the project of formal semantics is not be the only one, but it goes back to its founding fathers, and it is still probably the most common approach among its practitioners. The same is true about pragmatics: this take on pragmatics is not the only one, but it is one deeply rooted in the philosophical tradition of giving a common sense explanation of how speakers use language.

### 3. Data

Consider the following sentences:

1. The leaves are green
2. It is raining
3. The cat is on the mat

Now, consider the following scenario regarding (1) (due to Travis 1997, 89-90): Suppose Pia paints the leaves of a Japanese russet maple tree green for a photographic installation. Upon ending the job she might utter truly (1) while pointing to the leaves: (1) “The leaves are green”. Later a botanist friend seeking green leaves for green-leaves chemistry calls her. She might for all the paint, utter falsely (1) while pointing to the same leaves. Intuitively, the two utterances of (1) have different truth values; the first is true while the second is false.

Similar scenarios have been imagined for (2) and (3) Consider the following one due to Perry (1987/1993): “It is a rainy Saturday morning in Palo Alto. I have plans for tennis. But my younger son looks out the window and says (2) "It is raining." [...] Suppose, for example, that my son has just talked to my older son in Murdock on the telephone, [who tells him that it is sunny in Murdock] and in responding to my question [about Murdock he utters (2)]” (Perry 1987/199: 206-210). Intuitively the first utterance is true while the second is false<sup>37</sup>.

According to many authors neither (1) nor (2) and (3) are ambiguous, elliptical, contain vague or indexical expressions, and our intuitions about their truth values are not the result of what might be indirectly conveyed by their respective utterances (i.e. intuitions triggered by pragmatic implicatures)<sup>38</sup>.

---

37 I modified slightly both examples which are due to Travis and Perry respectively. I'll discuss (3) in more detail in the following chapter.

38 Although some authors reject this latter claim. See Cappelen and Lepore 2005, or Bach 2012. Skeptics about the possibility of formal semantics have put forward a plethora of cases similar with (1), (2) and (3). For examples see : Bach and Bezuidenhout 2002, Bezuidenhout 2002, Recanati 2001, 2004, 2007, Travis 1985, 1996, 1997, Searle 1978, 1992, Cohen 1985, 1986.

#### 4. Underdetermination

Why are these data problematic for standard semantic theories? According to skeptics standard semantic theories cannot generate correct truth-value predictions for (1), (2) and (3) in the imagined scenarios. And given that standard semantic theories aim to deliver specifications of truth conditions for all sentences of a language starting from the linguistic meaning of their simple expressions and a certain syntactic analysis of the language this is a severe failure. They fail to provide derivations of truth-conditions even for rather simple sentences like (1) to (3). According to skeptics whether (1) or (2) or (3) are true or false at the imagined scenarios depends not only on their respective linguistic meaning and how the world is, but on a multitude of potentially unrepeatable and formally intractable factors like the participants' immediate interests, purposes and concerns. In other words, the failure to give correct predictions of truth-values for utterances of (1) – (3) undermines the project of giving an analysis of natural languages that aims at predicting truth values of actual and possible utterances of their sentences<sup>39</sup>.

One of the main tenets of standard semantic theories is the claim that there is a determination relation between the linguistic meaning of sentences (their character) and their truth conditions, that for any context of utterance the linguistic meaning of a sentence determines how the world has to be in order for that sentence to be true at that context. Or as Searle puts it, the determination relation is the claim that “the meaning of a sentence determines a set of truth conditions; that is, it determines a set of conditions such that the literal utterance of the sentence to make a statement will be making a true statement if and only those conditions are satisfied” (Searle 1978: 208). This tenet follows directly from the fundamental insight that to know the meaning of a sentence is to know (or involves knowledge of) the conditions under which the sentence is true and, the conditions under which is false. And given that we cash out the notion of truth conditions is in terms of circumstances of evaluation and functions from circumstances to truth-values, we say that the linguistic meaning of a sentence determines a unique content with respect to any given context, and the content determines a unique truth-value with respect to any given

---

39 Searle 1978, Travis 1996, 1997, Ziff 1972, Bezuidenhout 2002, Recanati 2004

circumstance of evaluation.

We can summarize the main tenets that form the backbone of standard semantic theories in the following way:

(a) the linguistic meaning of sentences is determined by the linguistic meaning of their constituents and their syntactic structures (*compositionality of linguistic meaning*).

(b) for any context of utterance the linguistic meaning of a sentence  $S$  (i.e. what is linguistically encoded in  $S$ ) fully determines the content of  $S$  at that context, which in its turn, depending on how the world is, determines the truth value of  $S$  at that context (*the determination of truth-conditions by linguistic meaning*)<sup>40</sup>.

It is precisely the determination relation that skeptics believe that intuitions about truth-values of utterances of (1) - (3) falsify. They think that as far as the meanings of constituents of (1) and its syntax go whether we can predicate truly “green” of the leaves is an open matter: on some occasions we can on others we can’t. Or as Travis puts it “all that meaning fixes allows for words to state truth, but also falsehood, of given items in given conditions” (Travis 1996, 453). Allegedly the difference in truth-values is due to the fact that (1) has different truth-conditions at the two contexts of utterance: (1) is true in the photographer context iff the leaves appear green at the time of the utterance; while it is true in the botanist context iff the leaves are naturally green at the time of the utterance (i.e. if they are green in the way in which plants are green, namely green due to the presence of chlorophyl). Further on, it is argued by skeptics, that this difference in the truth-conditions of (1) is not determined by its linguistic meaning. Although the definite description “the leaves” is context-sensitive (it refers to the contextually salient leaves), the variation in truth-conditions of (1) is not due to its context-sensitivity since it refers to the same leaves at both contexts. Nor is the difference in truth-conditions of (1) due to the linguistic meaning of the predicate “is green”, since it refers to the same property at both contexts, namely the property of being green. Then the linguistic meaning of (1) should determine identical truth-conditions for (1) at the two contexts<sup>41</sup>. But (1) has different truth-conditions

---

40 Obviously the linguistic meaning of a sentence exhausts what is linguistically encoded in a sentence: what is encoded in the lexical meaning of its constituents and its syntax.

41 This is true only under the some further assumption: that "the leaves" and "is green" are the only constituents of (1) – there are no other constituents in the logical form of (1).

at the two contexts of utterance, or so it is argued. So the linguistic meaning of (1) fails to determine its truth conditions. Skeptics argue that the difference in the truth-conditions of (1) is due to the fact that the linguistic meaning of “green” does not by itself settle how an object must be in order to be correctly describable as green. On some occasions we can say of the given leaves that they are green (e.g. in the photographer context) but on other occasions we can't (e.g. in the botanist context). What determines whether an utterance of (1) is true in the photographer context and false in the botanist context is not just the linguistic meaning of (1) and the brute state of the leaves, but also the immediate interests and purposes of the participants in the conversation.

The argument for underdetermination can be summarized in the following way:

(A\*). The two utterances of (1) have different truth values.

(B\*) The difference in truth value is due to the fact that (1) expresses different contents at the two contexts of utterance.

(C\*) None of the constituents of (1) is context-sensitive so their semantic content does not vary across contexts of utterance.

(D\*) The linguistic meaning of (1) is determined compositionally from the linguistic meaning of its constituents.

From (B\*), (C\*) and (D\*) it follows:

(E\*) Failure of the determination relation: the linguistic meaning of (1) fails to determine the truth conditions of (1).

According to skeptics we need to look at a variety of factors, many of them context-specific and local (like the particular intentions, interests, and purposes) in order to determine the conditions under which a sentence is true or false relative to a context. Given that there are an indefinite number of potential contexts of utterance, there is no limit to the amount of pragmatic factors that can affect the truth-conditions of possible utterances of a sentence. Semanticists face, then, a serious obstacle: semantic theories must take into account an open ended number of factors and many of these factors do not seem amenable to a formal treatment as they are unsystematic, unrepeatable, hazy and imprecise<sup>42</sup>.

---

<sup>42</sup>This is the gist of Searle's (1978) argumentation, and of all *open-texture arguments*. I'll discuss them in more

The skeptic's conclusion is that the linguistic meaning of a sentence plays a role in fixing when the sentence would be true and when false, but not an exhaustive role. The linguistic meaning leaves room for variation in the truth conditions of a sentence from one context of utterance to another, variations that are not linguistically mandated. If we are to judge whether the utterance of a sentence is true we will have to look beyond the linguistic meaning of the sentence and how the world is. We must look at the immediate interests, particular purposes and the intentions that the conversational partners have at the context of utterance. Allegedly, we count an utterance of (1) or of (3) as being true only if it serves the purposes, interests and intentions that the conversational partners had at the time and place of the utterance. In fact, underdetermination goes against the fundamental semantic insight that to know whether a sentence is true (relative to a context) suffices that one knows the linguistic meaning of the sentence and how the world is at the time of context of utterance.

Then underdetermination is the claim that for some sentences and some contexts the linguistic meaning of the simple constituents of the sentence and its syntactic structure fail to determine, together with the (kaplanian) context, the truth conditions of that sentence relative to that context. This boils down to a claim that there is no systematic way for a semanticist to specify ahead what a sentence means relative to a context and, thus, no systematic way to predict its truth value relative to that context. Obviously this threatens the project of formal semantics. If the linguistic meaning of a sentence fails to determine under what conditions a sentence is true or false then semantic theories cannot generate, in a systematic manner, predictions about the truth values of that sentence.

According to many authors this phenomenon is not limited to (1) - (3) but it is pervasive, and some even go so far to claim that this sort of underdetermination affects virtually every English sentence. Skeptics take underdetermination to show that the entire project of formal semantics for natural languages is wrongheaded and doomed to fail (Travis 1996, 1997, Recanti 2004, Ziff 1972, Cohen 1986)<sup>43</sup>. But not all who accept

---

detail in the next chapter

43 Here is a clear formulation of this sentiment by Cohen (1986: 224) "we cannot construct a semantics for any natural language along the same lines as a semantics for a formal system of any currently familiar kind. Projects like Davidson's or Montague's cannot succeed."

underdetermination, take this nihilist stance: some take underdetermination to show that semantic theories need to adjust some of its main tenets. (Recanati 2010, Bezuidenhout 2002)

Obviously, in order for the underdetermination argument to have any force the data put forward by skeptics' scenarios must be semantically relevant: data that semantic theories must explain and must be accountable to. In other words, the intuition that the utterance of (1) in the photographer context is true and the intuition that its utterance in the botanist context is false must be the kind of intuitions that semantic theories must predict. But this didn't pass unchallenged.

Some theorists do not accept the intuitions presented above as being relevant for semantic theorizing. They reject the claim that the intuition that the two utterances have different truth-values is relevant for semantic theorizing (i.e that semantic theories must give predictions about them). They claim that intuitions put forward about the truth values of (1) in the imagined scenarios are inter-meshed with intuitions about what one does in uttering (1) given the intentions and interests of the conversational participants. These are, in fact, intuitions about whether what one asserts or conveys in uttering (1) is true, given the intentions and interests of conversational participants. Then the intuitions that one utterance of (1) is true while the other false are not to be explained by appeal to semantic properties of (1). For example, the intuition that (1) is false in the botanist context can easily be explained in the following way. We have the intuition that Pia's utterance of (1) is false because her answer the botanist leads him astray in believing wrongly that the leaves are green in the way in which leaves are normally green. In other words, the elicited intuitions about the truth-values of (1) are, in fact, intuitions about the truth-values of what is asserted or what is conveyed by its utterances. But obviously these are not the kind of intuitions relevant for semantic theorizing. In fact, some of the authors who take this line (e.g. Cappelen and Lepore 2005) go even further and claim that contrary to skeptics' intuitions (1) is true at both contexts. And not only it has the same truth-value at the two contexts, but it also has the same content, only that its corresponding utterances convey different things: the first utterance conveys the proposition that the leaves appear green,

while the second utterance conveys the proposition that the leaves are naturally green<sup>44</sup>.

Then one way to answer the skeptics' challenge is simply to circumvent it. It is not the job of semantic theories to account for the intuitions elicited by skeptics' scenarios simply because these are not the kind of intuitions that semantic theories must explain. The work of filtering out the data that semantic theories must explain and must be sensitive to from other features of linguistic communication leaves out of the scope of semantic theorizing the intuitions raised by these scenarios.

Obviously skeptics, and others who accept underdetermination, believe that the intuitions about the truth values of (1) elicited by the above scenarios are *bona fide* semantic data. Some formal semanticists working within the framework of this thesis concede this, and they believe that formal semantic theories can account for these data. And then, there are theorists who simply refuse to accept the intuitions about the truth values of (1) as being data that semantic theories must predict: allegedly what is communicated by the two utterances of (1) underlie our intuitions about the truth-values of (1), so these intuitions are not the kind of intuitions that semantic theories must generate predictions about. A great part of the debate regarding the skeptics' challenge concerns the status of these data: whether they fall in the ballpark of semantics (and thus it is semantic theories that must predict them) or whether they must be explained differently.

For the purposes of this thesis I will also concede that the data put forward by skeptics' scenarios are *bona fide* semantic data.

## **5. Underdetermination and Compositionality**

Underdetermination concerns the relation between the linguistic meaning of sentences and their truth conditions. But we might as well put the argument in favour of underdetermination in another way. As one that takes it to show that natural languages fail to satisfy strong compositionality: the content of sentences (relative to a context) is not

---

<sup>44</sup> Sainsbury (2002) Cappelen and Lepore (2005) argue that intuitions raised by Travis' scenarios are not relevant for semantic theorizing: they are not intuitions that semantic theories must predict and must be accountable to.



determined from the content of their constituents at that context and the way they are syntactically combined. Even if the constituents of (1) are supplied with constant contents at the two contexts of utterance, it is still not the case that (1) has unique or constant content at the two contexts of utterance. The argument goes along the following lines: (Travis 1994, Recanati 2004). The best explanation for the differences in truth value of (1) at the two contexts is that (1) has different contents at the two contexts. The manner in which the leaves are said to be green is part of the content expressed by (1) at each context, at the photographer context it expresses the content that the leaves appear green while at the botanist context it expresses the content that the leaves are naturally green. But, since the definite description refers to the same leaves in both contexts, then it has identical content at both contexts, and since the predicate is not context sensitive, but determines the same property (i.e. the property of being green) at any context of utterance, then it too has the same identical content at the two contexts of utterance. Then, the failure condition for strong compositionality obtains. Briefly, the argument is this:

- A. The two utterances of (1) have different truth values.
- B. The difference in truth value is due to the fact that the sentence expresses different contents at the two contexts of utterance.
- C. The semantic content of constituents of (1) does not vary across contexts of utterance.

From (B) and (C) it follows:

- D. Strong compositionality does not hold for English: there are two contexts and a complex expression, namely (1), such that the complex expression has different contents at the two contexts although the content of its constituents do not vary across the contexts of utterance. (Conclusion)

Of course, the argument has any bite only if (B) and (C) hold. We can block it either by denying (B) or by denying (C). So to go from intuitions about truth values of utterance of (1) to the conclusion that its semantics underdetermines its content two further steps are needed, namely (B) and (C). First one has to argue that (B) holds, that is that the content of (1) varies across contexts of utterance, that is to argue that the first utterance expresses the

content that the leaves appear green while the second that the leaves are naturally green. Secondly one has to argue that (C) holds, that is that the predicate of (1) is not context-sensitive (i.e. has a character that determine different properties at different contexts of utterance)<sup>45</sup>. In the next section I will use the framework introduced in chapter one to describe several ways in which one can account for the data, and to show how they relate with different premises of the argument.

Arguments in favor or against these two premises constitute the core of the debate concerning whether the data undermines the project of traditional semantics or not. I will present some of them further below only insofar as they are relevant to the discussion of how different positions that rest on accepting or denying B and C fare with respect to the different principles of compositionality. A further position accepts the entire argument and its conclusion and takes it to offer a good motivation for the claim that natural languages are best described as weakly compositional in the sense introduced in in chapter two. This corresponds to the less-skeptical position towards the project of traditional semantics: the project of giving a formal characterization of natural languages can be salvaged if we accept that semantic theories of natural languages are weakly compositional and accept pragmatic effects on content.

The underdetermination thesis claims that for some sentences of natural languages (possibly for all sentences) a theorist cannot specify ahead all of the possible truth-conditions of those sentences. This is a claim about the relation between the linguistic meaning of sentences and their truth conditions and ultimately about whether semantic

---

45 In fact this is but one of the things that a supporter of underdetermination has to show. She also has to argue against ambiguity (if “green” is an ambiguous term this can explain the difference in truth value), vagueness (being a color term, “green” is vague, so she has to argue that the difference in truth values of (1) is not due to vagueness of “green” but its orthogonal on that), or gricean domestications ( the claim that (a) the two utterances of (1) express the same content, but convey different things, where the second utterance conveys the proposition that the leaves are naturally green, which is false and (b) that the intuition about the truth value of (1) is in fact an intuition about the truth value of what is conveyed by that particular utterance of (1)). She would also have to argue that the intuitions about utterances of (1) hold water. One can deny the data and claim that both utterances of (1) are true: one can either claim that (1) is an example of unspecific meaning and that both utterances are true but the second one is misleading (Sainsbury 2002, Berg 2002) or claim that (1) has a constant content across *all* contexts of utterances, and that the intuitions about the two utterances are the result of what is communicated by them (Cappelen and Lepore 2005). Some defenders of underdetermination (e.g. Travis 1997, Recanati 2004 and Bezuidenhout 2002) consider all these alternatives and put forward arguments against them. At least for some of the examples given, the issue is not settled yet. I will not review the arguments that purport to show that (1) is or isn’t any of the above, unless the points that I will discuss below hinge on them.

theories can generate, in a systematic way, correct truth-value predictions about possible and actual utterances of those sentences.

From the argument given above it follows that if underdetermination is true of a language (or a fragment of it) then also failure of strong compositionality is true of that language. In other words, failure of the determination relation, brings about failure of strong compositionality. If this is true, so is its contrapositive: a semantics that satisfies strong compositionality is a semantics that satisfies the determination relation.

This brings us to an important aspect of the relation between underdetermination and strong compositionality. One way to resist the argument from underdetermination and, thus, one way to resist skepticism, is to give a descriptively adequate semantics that satisfies strong compositionality. For, example, if a theorist gives a strongly compositional semantics for a fragment of English that includes (1) to (3), one will, in fact, show that the underdetermination arguments do not hold water.

But notice that although failure of the determination relation entails failure of strong compositionality, its converse is not true: failure of strong compositionality does not entail failure of the determination relation as shown by semantics that contain context-shifting operators (like the AtC2 operator discussed in chapter two). Such a semantics will fail to be strongly compositional but it can satisfy character compositionality *and* the determination relation, in the sense that the linguistic meaning it assigns to sentences will determine, for any given context, their truth-conditions at that context.

The claim of underdetermination is that the linguistic meaning of sentences fail to determine their content at possible contexts of utterance. This supports the skeptical claim that there we can't give a theory that models the repeatable and conventional aspects of meaning (i.e. linguistic meaning) by giving predictions about the correct uses of sentences (i.e. truth-value predictions). Accepting underdetermination amounts to the upholding the claim that there is no way for the theorist to specify in a theoretically useful way the linguistic meaning for expressions of a language. Even if the theorist manages to specify some conventional aspects of meaning associated with sentences, this still fails short of determining their truth-conditions relative to contexts of use.

## **6. How to Account for the Data: A Brief Geography of Positions**

There are three general reactions to the underdetermination thesis. A radical one accepts underdetermination and takes it to show that formal semantics is not possible. The opposite reaction is occupied by a host of positions that defend (the project of) formal semantics that reject the underdetermination claim. A commonly employed strategy by theorists who defend the project of formal semantics and who accept that the intuitions about (1) to (3) raised by skeptic scenarios are relevant for semantic theorizing is to argue that a strongly compositional semantics for fragments of English containing (1) – (3) is possible. This strategy is followed for the reasons given above, although they are never made explicit. Finally, there is a middle-ground type of position that accepts the underdetermination argument but claims that a large part of the project of formal semantics can be salvaged if we are ready to accept some substantial changes in our semantic theories.

Besides the skeptics who believe that the data cannot be accounted by a systematic theory many other authors believe otherwise. There are several accounts that compete in explaining the data. The general framework allows us to ask several questions with respect to the data, and each possible answer delineates such an account:

Is the difference in truth value of utterances of (1) due to a difference in content?

No. The same content

Yes. Different contents

Does the semantics of (1) determine the content of (1) relative to a context?

No: the content of (1) is underdetermined by the semantics of (1)

Yes: the semantics of (1) determines the content of (1)

Variables in syntax

Indexicals

A first account of the data starts by denying premise (B) of the underdetermination argument; that is to say that it answers ‘no’ to the first question: the difference in truth value of utterances of (1) is *not* because (1) expresses different contents at the two contexts of utterance. I will call this position, *relativism*. An intuitive way to look at the data, is to say the two utterances of (1) have different truth values because the manner in which the leaves are green matches the interests and purposes of the painter, but not those of the botanist. We can say that, in a sense, the leaves are green by the standards of the painter, but not by the standards of the botanist. The semanticist working in the framework of this thesis has two options for modeling the role that the manner of being colored plays in determining the truth values of (1). According to relativism, (1) expresses the same content at the two contexts, and the content expressed is unspecific with respect to the manner of being colored. The difference in truth-values comes from the fact that the content expressed gets evaluated at different circumstances of evaluation, and the manner of being colored is an index of the circumstances of evaluation. The value of the manner-of-being-coloured index at which we evaluate the content expressed by (1) is determined by what is salient to the participants in the conversation<sup>46</sup>. Thus the content expressed (1) at any context is the relative proposition that the leaves are green. This relative proposition can be represented as a function from triples consisting of world, time, manners-of-being-colored to truth values, such that it returns true relative to a triple  $\langle w, t, m \rangle$  just if the leaves are green according to the manner  $m$ , at the time  $t$  in world  $w$ . Since the interests and purposes of the photographer differ from those of the botanist, technically speaking they determine different manner-of-being-colored indices. Then, (1) has different truth values at the two contexts of utterance, because the content it expresses at the two contexts, gets evaluated at circumstances that differ with respect to the manner of being colored index. Such an account is compatible with strong compositionality.<sup>47</sup>

On the other hand if we accept that the difference in truth value of the two utterances of (1) is because of the differences in the content of (1), then there are several ways to account for the data. One line starts by denying premise (C ) of the

---

46 Some theorists might say that the manner-of-being colored index is determined by the intentions and purposes of the conversational partners.

47 Predelli (2005) defends such an account.

underdetermination argument by claiming that the content of some constituents of (1) vary across contexts of utterance. Further along this line one can argue that it is the content of constituents articulated in the surface form of (1) that vary their content. If so, it might be reasonable to treat them in the same way in which we treat indexicals like “I” or “that”: surprisingly, “green” turns out to be an indexical. I will call this position *indexicalism*. Depending on how one treats indexicals this position can also be compatible with strong compositionality. According to indexicalism the fact that a sentence like (1) receives different truth valuations relative to different contexts is no different than the fact that a sentence like “I am hungry” receives different truth valuations relative to different contents. Indexicalism follows a second option that the semanticist working in the framework of this thesis has in modeling the role that the manner of being colored plays in determining the truth value of (1). Again, intuitively, Pia, the photographer, is interested in leaves that appear green, while her botanist friend is interested, for bio-chemistry experiments, in leaves that are naturally green. So, in a sense, the leaves count as green according to Pia's standard, but not according to the botanist's standards. The indexicalist proposal is to take the manner-of-being colored as a contextual parameter determined by the interests and purposes of participants, and take linguistic meaning of “is green” to be sensitive to this parameter<sup>48</sup>. The linguistic meaning of “is green” can, then, be modeled as a non-constant function from contexts to properties, such that at contexts that differ with respect to the manner-of-being-colored parameter it returns different properties<sup>49</sup>.

There is yet another position that denies premise (C) of the underdetermination argument and claims that the variation in content of (1) across contexts is due to there being constituents of (1) that have different contents at different contexts. The difference with indexicalism lies in the fact that according to this account what varies across contexts is the semantic values of semantically relevant constituents present in the logical form of (1) but

---

48 Some formal semanticists might find direct appeal to intentions unpalatable. Then maybe another way to articulate indexicalism is by saying that "is green" is sensitive to the conversational standards of the context, namely the contextually *most salient* standard of color. Something is made salient by being directly available to the conversational participants, or by what has been previously said in the conversation, by what is commonly known between the conversational partners, and so on. Obviously salience is not a wholly objective notion, since what is salient is always salient to someone, but the hope is that it can be specified independently of speakers' intentions. A similar move is available to the relativist if she wants to avoid speaking of circumstantial indexes being determined by speakers' interests.

49 Rothschild and Segal (2009) defend such an account.

not articulated in its surface form. According to this position, although it is not immediate obvious, all aspects of truth conditions of utterances of (1) are traceable to its formal features. The claim is that the logical form of (1) differs from its surface forms in that it contains constituents that are not articulated in the surface form. Let's call this the *hidden indexicals-account*. This position is also compatible with strong compositionality. There are several possible ways to implement this idea, depending on what kind of elements get posited in the logical form of (1) and where in the logical form they get posited. One proposal is that "green" co-habits its syntactic node with two variables: a variable that takes as value a comparison class provided by the context, and a variable that takes as value a certain part of the object, which is also provided by the context. In order to determine the truth-conditions of (1) relative to a context of utterance, the variables in its logical form must be assigned a value. The content of (1) varies across contexts of utterance precisely because the value of its variables varies across contexts of utterance<sup>50</sup>.

These three positions briefly sketched here, both accept the data put forward by skeptics as *bona fide* semantic data, and are all squarely within the framework introduced in chapter one. Finally, there is another position which accepts the data, but departs from some of the important tenets of the framework introduced in chapter one. This position accepts that the difference in truth value of the two utterances of (1) is due to differences in the content of (1), but denies that the literal meaning of (1) determines its content relative to a context. This is the line taken by *truth conditional pragmatics*. TCP fails to satisfy strong compositionality but it *can* satisfy weak compositionality. According to this position we can accept underdetermination of truth-conditions by linguistic meaning, but can still hope for a systematic account of the semantics of natural languages. TCP aims to explain the aspects of linguistic competence with a language that traditional semantic theories hope to explain. And it aims to explain them also by generating systematic truth-value predictions of possible and actual utterances of that language. Some of these aspects include explaining intuitions about truth-conditions, explaining the fact that linguistic competence is infinite in its scope, and that competent speakers have the ability to understand novel sentences. There are two changes that we are asked to accept: that pragmatic factors play a role in

---

50 Szabó (2001) and Stanley (2002) defend such an account.



determining the truth conditions of sentences. And secondly, that we should give up the requirement that the assignment of content satisfies strong compositionality and replace it with a weaker one that the assignment of content satisfies weak compositionality. In other words, we are told we can still do what formal semantic theories hoped to do, only that we should do it by other means: by accepting that pragmatics intrudes in the determination of truth-conditions.

## 7. Truth Conditional Pragmatics

According to skeptics underdetermination of truth-conditions by linguistic meaning endangers the project of formal semantics because it undermines the capacity of theories to make truth-value predictions. Now, truth conditional pragmatists accept underdetermination but claim that a systematic account of the meaning-properties of natural languages can still be given, if we accept (a) that such an account is weakly compositional as defined in chapter two and (b) that truth-conditions are partly determined by pragmatic factors (i.e. pragmatic intrusion)<sup>51</sup>. Truth-conditional pragmatics (henceforth TCP) is put forward as an alternative to the standard semantic theories done within the framework introduced in chapter one.

Acceptance of pragmatic intrusion offers not only a new take on the relation between semantics and pragmatics but it also brings important differences between TCP and traditional semantic theories. It offers a new perspective on what pragmatic factors do, but also on what semantic content is, and the role it plays in semantic theories. To understand the claims of TCP better, consider, again, the traditional picture that TCP seeks to replace.

Theories done in the traditional framework make several important claims about semantic content. A first one is the determination relation: the semantic content of a

---

51 Because such theories still aim to match sentences with their truth-conditions (relative to contexts) I will continue to call them semantic theories, even though they incorporate pragmatic explanations.

sentence is determined by its linguistic meaning and the context of utterance. A second one is that semantic content is propositional in nature in the sense that it can be represented as a function such that given a circumstance it returns a truth value. A third claim is that theorists have a handle on semantic content of sentences thanks to intuitions about their truth conditions (i.e. intuitions about truth values of their actual and possible utterances). Given that sometimes what is conveyed by an utterance of a sentence is not what the sentence literally means, semantic theories come supplemented with an account of what one conveys by an utterance over and above what the uttered sentence means at that context of utterance. Supplemented in this way standard semantic theories are committed to two further claims. One is that pragmatic factors do not affect the semantic content of a sentence at a context of utterance but kick in only in as much as to determine what is conveyed by an utterance (when what is conveyed by that utterance is not exhausted by the semantic content of the uttered sentence). The second claim is that what a sentence means relative to a context (its semantic content) is the input for the derivation of what an utterance of that sentence conveys (i.e. to implicature derivation). The general picture then is the following: the content of a sentence *S* at a context *C* is determined by the linguistic meaning of *S* and the context of utterance. What is conveyed by an utterance of *S* on the other hand is determined by the content of *S* at *C* and pragmatic considerations regarding the speaker's cooperativeness and the appropriateness of uttering *S* on that occasion.

TCP rejects this picture almost entirely. It claims that pragmatic factors have a role to play not only in determining what is conveyed but also in determining the semantic content of sentences. In other words, pragmatic factors play a role not only secondary to the determination of semantic content but they are partly responsible for determining the semantic content itself. This is a consequence of the rejection of the determination thesis. TCP accepts, though, the other claims about semantic content: that they are propositional in nature (return a truth value given a circumstance), that they provide the input to implicature derivation, and that thanks to intuitions about truth conditions theorists have a reliable grasp of them.

TCP theorists argue that the content of each of (1) to (3) relative to a context goes beyond what is encoded in the sentence itself and is partially determined by free pragmatic

factors “which are not triggered by an expression in the [uttered] sentence but take place for purely pragmatic reasons – to make sense of what the speaker is saying”(Recanati 2011, 6). Here is how Recanati summarizes the main idea behind truth conditional pragmatics:

“truth conditional pragmatics is the view that the effects of context on the [semantic] content need not be traceable to the linguistic material in the uttered sentence. Some effects of context on [semantic] content are due to the linguistic material (e.g. the context sensitive words or morphemes which trigger the search for contextual values), but others result from “top down” pragmatic processes that take place not because the linguistic material demands it, but because utterance’s content is not faithfully or wholly encoded in the uttered sentence, whose meaning requires adjustment or elaboration in order to determine an admissible content for the speaker’s utterance”. (Recanati 2011, 127)

What are these pragmatic factors? They can't be those relevant for deriving what the speaker conveys by uttering an sentence at a context, for such factors require first that what the sentence means at the context is already determined. Implicature derivation takes semantic content as input. On the other hand such pragmatic factors have semantic content as output.

Here is my reconstruction, which I believe is a plausible formulation on behalf of the TCP theorist of how we should understand the difference between pragmatic and semantic factors that affect truth-conditions. Something counts as pragmatic (as opposed to semantic) if it is not linguistically mandated: not required by what is linguistically encoded in a sentence<sup>52</sup>. The mark of the pragmatics is optionality, while the mark of the semantics is being linguistically mandated. A factor  $p$  is mandatory with respect to a sentence  $S$  iff there is no context of utterance  $C$  such that  $S$  can be evaluated for truth-value at  $C$  in the absence of  $p$ . A factor  $p$  is optional with respect to a sentence  $S$  iff there is at least one context of utterance  $C$  such that  $p$  is relevant for the truth-evaluation of  $S$  at  $C$  and there is a context of utterance  $C^*$  such  $S$  can be evaluated for truth value at  $C^*$  in the absence to  $p$ . This way of dividing semantics and pragmatics does not overlap with divide between the determination of semantic content and implicature derivation, the but is orthogonal to it<sup>53</sup>.

---

52 In a sense this comes down to a stipulation of the use of “pragmatic”.

53 Theorists that accept pragmatic intrusion must make two orthogonal distinctions. First they need to distinguish between semantic factors and pragmatic factors that determine the content of an expression. And when it comes to pragmatic factors they need to distinguish those pragmatic factors that play a role in determining the content of an expression, from pragmatic factors that determine what is conveyed by uttering that expression at a given context.

An example will help shed light on the difference between the optional and the mandatory.

In order to determine whether (I) “I’m hungry” is true or not relative to a context we need to determine who the speaker is and whether the speaker is hungry at the time of the utterance. In other words, there is no context of utterance C such that we can evaluate (i) for truth at C unless the referent of the first person pronoun “I” is fixed. In this respect fixing the reference of “I” relative to a context is mandatory – without it (i) is not truth evaluable. Indexical saturation is mandatory also in another sense: it is *mandated* by the linguistic meaning. Presumably, not only fixing the reference of indexicals is mandatory but also assigning values to variables in a sentence, both articulated in the surface form and unarticulated in the surface form (but present in the logical form). For example the genitive phrase “John’s team” in (ii) “John’s team lost” can be interpreted either as referring to the team that John plays for, or that John owns, or that John cheers for, the team that he betted on, the team that he coaches, and so on, allegedly with virtually no limits to the relevant relation. In order to truth-evaluate an utterance of (ii) it is necessary to determine what the relevant relation between John and the team is, for it might be that the team that John owns lost while the time that he coaches won. We can analyze “John’s team” as containing a free variable R, where R stands for the relation between John and the team: *the team that bears relation R with John*<sup>54</sup>. The genitive phrase “John’s team” expresses a content only if the free variable R is given a value, and the value is always contextually determined by “looking for an appropriately salient and relevant relation in the linguistic or non-linguistic context” (Partee 2004: 118). Indexical resolution and assigning values to variables are *mandatory*, in the sense that no sentence containing indexicals or variables can be evaluated for truth unless indexical resolution and variable assignment take place.

According to TCP not all factors that determine the truth-conditions of an expression relative to a context are linguistically mandated. The ways in which the leaves must be green for (1) to count as true or false is not encoded in its linguistic meaning. The intentions and interests of conversational partners might be relevant for the truth-evaluation of (1) relative to some contexts, is not linguistically encoded in (1). Or consider (2) “It is

---

54 Admittedly this is very rough. There are different ways to make this more precise. For example Partee analyzes “John’s team” as a definite description containing a free variable R: the  $x$  such that team ( $x$ ) and R (John,  $x$ ). (Partee 2004:118-119)

raining” and Perry’s scenario. In order to evaluate whether its first utterance is true we need to check whether it is raining in Palo Alto. To evaluate whether its second utterance is true we need to check whether it is raining in Murdock. At least, *prima facie*, it seems that we can evaluate an utterance of (2) as true or false only relative to a location<sup>55</sup>. Now, if location is mandatory for the truth evaluation of (2) and location is part of the content of (2) relative to a context, then determining the location in (2) is on a par with saturating the indexical in (i). If so, then there should be no context C such that an utterance of (2) at C is truth evaluable although no location is provided. But Recanati (2002) argues that we can imagine a possible context of use C, such that (2) is truth evaluable at C although no location is provided at C: we can evaluate (2) for truth, in the absence of a location. Then, the provision of a location in the Murdock-context and in the Palo Alto-context is not the result of what is linguistically encoded in (2) but the result of pragmatic reasoning.

I will not evaluate directly the arguments that try to show that there are optional factors (i.e. pragmatic factors) that are truth-conditionally relevant. Arguments that try to show this, or the opposite, constitute the core debate concerning context-sensitivity and the viability of formal semantics. In fact, most authors approach this debate by putting forward such arguments: arguing either for or against pragmatic intrusion. Rather, for the sake of the discussion, I’ll concede to the truth-conditional pragmatist that pragmatic factors do affect truth-conditions, and ask what does this mean for compositionality.

---

55 The framework offers two options to account for this. One is to claim that the location is part of the content of (2) relative to a context of utterance. Thus (2) at the first context of utterance expresses the content that it is raining in Palo Alto (which can be represented as a function from world time pairs  $\langle w, t \rangle$  to truth values, such that it yields truth iff it rains in Palo Alto in  $w$  at  $t$ ) while at the second context of utterance it expresses the content that it is raining in Murdock (which can be represented as a function from world time pairs  $\langle w, t \rangle$  to truth values, such that it yields truth iff it rains in Murdock in  $w$  at  $t$ ). The other option that the framework gives us is to claim that the content of (2) at the two contexts is the same and to represent it as a function from world-time-location triples  $\langle w, t, l \rangle$  to truth-values, such that it yields true if it is raining in  $w$ , at  $t$ , in  $l$ .

## 8. Pragmatic Intrusion and Weak Compositionality: The Model

In this section I will give a brief summary of how TCP formally accommodates pragmatic operations in the theoretical machinery, and the role of weak compositionality in this. One important aspect of pragmatic intrusion is that it is context-specific in the sense that *what* pragmatic factors determine the content of a sentence like (1) varies across contexts of utterance and is not fixed in the way in which the contextual parameters relevant for determining the content of indexicals like “I” or “now” are fixed. For example, whether the leaves counts as being green depends on whether we are in a domestic day-to-day environment, or in a biology lab, or on a photographic set, and so on. A second important feature of pragmatic intrusion is that it affects both simple and complex expressions alike.

Here is a brief reconstruction of the truth-conditional pragmatic model of natural languages. For simplicity I'll adopt Recanati's term and call any pragmatic operation that is truth-conditional relevant, “modulation”. An informal presentation of how the truth-conditions of (1) are determined according to TCP will offer an intuitive grasp of the idea. According to TCP the linguistic meaning of each simple constituent expression of (1) determines together with the context of utterance the content of that expression at that context. The result of combining these contents according to the syntactic structure of (1) is the proposition that the leaves are green (i.e. a function that returns truth iff the leaves are green in one way or another). According to truth-conditional pragmatists this proposition can't be the content of (1) in the botanist-context. The content of (1) at the botanist-context are delivered by modulation functions that take the proposition that the leaves are green and yield the contextually appropriate proposition that the leaves are naturally green. What results from combining the semantic content of the parts is only an intermediate stop in the overall process of determining the truth-conditions of (1). Since for any given expression the pragmatic factors relevant for determining what that expression means relative to a context vary with the context, at different contexts of utterance different modulation functions determine the content of that expression at that context. For each expression-

context pair there is a particular modulation function that determines the correct content of  $e$  at  $C$ : the contextually salient/relevant/apropriate function for the interpretation of  $e$  at  $C$ . Formaly we can define a general modulation function  $mod$  such that for any expression-context pair  $\langle e, C \rangle$ ,  $mod$  yields the contextually relevant/salient/appropriate modulation function which determines the content of  $e$  at  $C$ :  $mod(e, C)$ . In other words the general modulation function  $mod$  is a (partial) function defined over the cartesian product of the set of contexts and set of expressions with values in the set of contextually specific modulation functions. In their turn each contextually specific modulation function takes the content (intension) determined by the linguistic meaning of  $e$  together with the context of utterance  $C$ , and returns a content (intension). (Contextually specific modulation functions - e.g.  $mod(e, C)$  - are functions from intensions to intensions). I will write the value of a particular modulation function as:  $mod(e, C)\mathbf{I}(e, C)$ . These values are the building blocks which the composition function takes to deliver the content of complex expressions. Then, as a first aproximation we should say that the content of a complex expression is a function of the modulated contents of its parts and the way they syntactically combined:

$$\mathbf{I}(\alpha(e_i, e_j), C) = f(\alpha, mod(e_i, C)\mathbf{I}(e_i, C), mod(e_j, C)\mathbf{I}(e_j, C))$$

But this is not the whole story. Pragmatic intrusion affects both simple and complex expressions which is to say that modulation can operate on complex expressions without operating on their simple constituents<sup>56</sup>. The content of a sentence, for example, can itself be the result of a contextually salient modulation function:

$$\text{Mod}(\mathbf{I}(\alpha(e_i, e_j), C)) = mod(\alpha(e_i, e_j), C)(f(\alpha, mod(e_i, C)\mathbf{I}(e_i, C), mod(e_j, C)\mathbf{I}(e_j, C)))$$

Notice that this is equivalent with the following formulation:

$$\text{Mod}(\mathbf{I}(\alpha(e_i, e_j), C)) = f \circ mod(\alpha(e_i, e_j), C)(\alpha, mod(e_i, C)\mathbf{I}(e_i, C), mod(e_j, C)\mathbf{I}(e_j, C))$$

where  $f \circ mod(\alpha(e_i, e_j))$  is a complex function obtained by composing the composition function  $f$  and the contextually relevant function  $mod(\alpha(e_i, e_j)C)$ .

---

56 Here is how Recanati puts it “The distinction between [semantic] content and modulated [content] applies to complex expressions as well as to simple ones: the [semantic] content of a complex expression is a function of the modulated [contents] of its parts and the modulated [content] of the expression results from modulating the [semantic] content thus determined.” (Recanati: 2010, 45) And he continues: “the modulated [content] of the complex is not a function of the modulated [contents] of its parts (and the way they are put together): for a given complex with a given content (determined by the modulated meanings of its parts and the way they are put together) can still be modulated in different ways.

Obviously, an account in which the content of complexes is determined by a context-specific modulation content fails to be strongly compositional. In such an account it is possible that the modulated content of the complex varies across contexts of utterance although the content of their constituents is stable across the very same contexts of utterance because any given complex expression can be modulated in different ways at different contexts. For any complex expression  $\alpha(ei, ej)$ , and any two contexts C1 and C2, it is possible that the modulation function relevant at the two context differs:  $mod(\alpha(ei, ej), C1) \neq mod(\alpha(ei, ej), C2)$

But this is compatible with, or is allowed in a account that satisfies weak compositionality. Weak compositionality allows that the content of a complex expression varies across contexts of utterance, but that the content of its constituents remain stable across the very same contexts of utterance. Moreover, the very idea of weak compositionality, that of letting the context as an extra argument of the composition function, makes room for pragmatic factors that determine the content in a way that is not mandated linguistically.

## 9. The Wrong Type of Arguments from Context Sensitivity

There is another argument involving context-sensitivity that allegedly shows that natural languages are not compositional. I think this later argument is wholly misguided and it is not a good way to argue that natural languages are not compositional. Nevertheless the argument is a particular instance of a more general type of arguments against compositionality, and it spawned a considerable literature so, I'll mention it here in order to set it aside<sup>57</sup>.

We are asked to consider sentences like the following:

(5) John likes red watermelons.

---

<sup>57</sup> See Lahav 1989, Reimer 2002, Recanati 1995, 2004, Szabó 2001, Janssen 1997, Fodor and Lepore 1992. See Pelletier 1994/2004 for a criticism of this kind of argument.



- (6) John doesn't like red apples.
- (7) John owns a red house
- (8) John didn't see the red traffic light
- (9) John needs red ink

Intuitively the meaning of (6) is that John likes apples that are red skinned, while the meaning of (5) is that John likes watermelons with red pulp. Some authors<sup>58</sup> claim that “nothing in the [meaning] of “red” or in that of “apple” suggests that the color of apple is identified with the color of its skin” as nothing in the meaning of “red” and that of “watermelon” suggest that the colour of watermelons should be identified with the colour of their pulp. (Reimer 2002:187). They take this to show that there are things in the meaning of the complex expression that just are not in the meanings of the parts and the syntax. Saying that the meaning the complex expression goes beyond the meanings of the parts and its syntax presumably is one way of saying that the former is not a function of the latter but that it requires something more.

A short reply is the following: this does not constitute an argument against any of the principles of compositionality discussed above since the failure condition of any principles is not met. More exactly, (5) and (6) do not show that English doesn't satisfy strong compositionality or that it doesn't satisfy character compositionality. This is so because neither the failure condition of strong compositionality is met, nor the failure condition of character compositionality is met. To show that English doesn't satisfy strong compositionality one needs to show that two utterances of (6) have different contents although its constituents have stable contents across the two contexts of utterance. To show that English is not character compositional one would have to show that (6) and a sentence obtained by replacing within (6) one of its constituents with its synonym have different characters. This is not what (5) and (6) show.

Other authors have a different take on these examples. They start from the observation that in order for (5) to be true, some watermelons must be in the set of things that are red on the inside, for (7) to be true John's house must be in the set of things that are

---

58 Lahav 1989: 401-403, Sainsbury 2002: 187, Reimer 2002:195, Recanati 2010: 29-32

red on the outside, and for (8) to be true, the traffic light must be in the set of things that emit red light, and for (9) to be true, ink must be in the set of things that appear green (when dried on paper surface). They claim that these examples show that “red” changes meaning when it is part of different larger expressions, and that its meaning is partly determined by the meanings of the other constituents of the larger expression. The adjective “red” has a different meaning depending on the noun it modifies, and in each case it refers to a different set of things. The meaning of “red” changes depending on the meaning of co-occurring expressions. Furthermore they claim that if the meaning of an expression is determined by the meaning of its co-occurring expressions then compositionality is not satisfied. Presumably the idea is that compositionality requires or entails that simple expressions have meaning independently of their occurrences in larger expressions and that they do not change their meaning when they become part of (different) larger expressions<sup>59</sup>.

Again, the short reply is the same as above: the failure condition of any of the principles of compositionality is not met, so there is no failure of compositionality. In fact, a semantics that assigns meanings to simple expressions as a function of the meanings of other co-occurring constituents is allowed by all the above principles of compositionality. It doesn't follow that if a semantics **I** assigns content to simple expressions as a function not only of the context of utterance but also of the content that other co-occurring expressions have at that context, then **I** is a semantics in which for some expressions  $ei$ ,  $ej$ , syntactic operation  $\alpha$  and some contexts  $C_1$ ,  $C_2$ ,  $\mathbf{I}(ei, C_1) = \mathbf{I}(ei, C_2)$  and  $\mathbf{I}(ej, C_1) = \mathbf{I}(ej, C_2)$  but  $\mathbf{I}(\alpha(ei, ej), C_1) \neq \mathbf{I}(\alpha(ei, ej), C_2)$ <sup>60</sup>.

So for the rest of the thesis I will not be concerned with any of this kind of arguments.

---

<sup>59</sup>See Lahav 1989: 402-404, and Recanati 1995 and 2004: 135-136 for such a claim. Fodor and Lepore (1992) too make a big deal about the claim that if the meaning of an expression depends on the meaning of other co-occurring expressions then compositionality fails to obtain. But they use it to prove the opposite. They think that natural languages are un-doubtedly compositional; hence the meaning of a simple expression cannot depend on the meanings of other co-occurring expressions.

<sup>60</sup>Pagin (1997) shows that a semantics where the meanings of simple expressions are dependent on the meanings of other simple expressions to the extent that the meaning of any simple expression depends on the meaning of every other simple expression is compatible with compositionality. We can think of a semantics where the meaning of simple expressions is a function of the meanings of other co-occurring simple expressions as a special case of this other more radically holistic semantics where the meanings of all simple expressions are inter-dependent.

## **10. Summing up**

In this chapter I presented the main challenge against the very possibility of formal semantics. According to skeptics for virtually any natural language sentence its truth-conditions are severely underdetermined by linguistic meaning. In this chapter I also explained how we should understand the idea of underdetermination and its relation with various principles of compositionality. The determination of truth-conditions of virtually any sentence is susceptible to the influence of an open-ended number of contextual factors. This endangers the capacity of formal theories to derive truth-conditions for sentences at context, and a fortiori their capacity to make truth-value predictions. One consequence of underdetermination is failure of strong compositionality: if sentences of a natural language exhibit underdetermination, then that language cannot receive a strongly compositional semantics. Finally, I ended with detailed presentation of truth-conditional pragmatics and its relation with strong and weak compositionality.



## **CHAPTER 4: Weak Compositionality and Systematicity.**

### **1. Introduction: Systematicity and Derivation of Truth-Conditions**

The general picture we get from the previous chapter is the following: according to skeptics the project of formal semantics is wrongheaded: because the truth conditions of virtually any natural language sentence vary freely with an open ended number of pragmatic factors it is impossible to give a systematic semantics to natural languages. In other words, because pragmatics intrudes so thoroughly in determining truth conditions, a semanticist cannot specify ahead what an expression means relative to a context and thus semantic theories cannot give correct truth-value predictions. The thesis of underdetermination threatens the project of formal semantics: if the linguistic meaning of a sentence plus (kaplanian) context fails to determine how the world has to be in order for that sentence to be true in that context, then there is no systematic way for the semanticist to determine in advance what its truth conditions relative to a context are and there is no way to give predictions about the truth values of utterances of that sentence. In that case the entire project of formal semantics implodes.

Supporters of weak compositionality argue that we should not despair: we can accept that pragmatic factors partly determine truth conditions and, nevertheless, give a systematic semantics to that language if we make use of weak compositionality. In other words, weak compositionality offers the best of both worlds: it allows both for pragmatic intrusion in the determination of truth conditions and for a systematic semantics.

I will argue that contrary to what supporters of weak compositionality claim, weak compositionality is not going to help with systematicity. A weakly compositional semantics that allows for pragmatic intrusion can still be un-systematic in the sense that it cannot systematically derive truth-conditions therefore cannot generate systematic truth-value predictions.

But first, I'll say a few words about what is it for a semantics to be systematic and what is it to derive truth-conditions in a systematic manner. Speakers' ability to make judgments about the correct use of expressions, including their ability to make judgments about how the world has to be in order for a sentence to be true, constitute their linguistic competence with a language. Given that semantic theories are interested in modeling this linguistic competence an important criterion of empirical adequacy for them is their ability to predict correctly speakers' judgments (intuitions) about correct use. And an important portion of these judgments are judgments about the conditions under which particular sentences are true: judgments on whether actual or possible utterances of a given sentence are true or false given a certain state of the world. Any semantic theory worth its name must, then, pair sentences with their truth-conditions. Once it does this, it can predict for any sentence of the language and for any possible state of affairs if that sentence is true at that state of affairs<sup>61</sup>. It goes without saying that the pairing of sentences with their truth-conditions must be done in a systematic manner.

And again, in order to make predictions about natural language sentences the theorist must first decide what intuitions about truth-values of natural language sentences are relevant for semantic theorizing, that is, what intuitions constitute the data that semantic theories must explain. Considerations that distinguish between what is conveyed by an utterance and what the uttered sentence means at that context help filter those intuitions that are to be explained by semantic theories from those that are to be explained differently. Secondly, the theorist must decide how predictions of truth values for sentences of the model language relate to speakers' intuitions about truth values of utterances of natural language sentences. This is taken care of by considerations that determine whether the model language (the language as described by the theory) is *the actual language* of a given group or population.

But how do semantic theories derive truth-conditions in a systematic manner? More precisely, how do semantic theories pair sentences of the model language with their truth

---

61 This is an idealization because of indeterminacy. For some sentences and some states of affairs speakers don't have determinate intuitions on whether that sentence is true or false at that state of affairs. For some sentences and some states of affairs we wouldn't know what to say with respect to the truth-value of that sentence at that state of affairs.

conditions and subsequently make predictions about their truth-values at given states of affairs? This is taken care of by the inner workings of the semantic machinery.

This is what I'll focus on in this chapter: the inner workings of semantic theories and how weak compositionality fits in. I will argue that if we accept that pragmatics intrudes in the determination of truth-conditions this endangers the capability of theories to derive truth-conditions in a systematic manner. Furthermore, I will argue that even though weak compositionality brings about *some form of systematicity* it is of no help when it comes to systematic derivations of truth-conditions.

For the sake of discussion consider, again, a broad sketch of how a semantics  $\mathbf{I}^*$  (that assigns linguistic meaning as semantic values) derives truth-conditions and thus how it make truth-value predictions. How such a semantic theory assign, in a systematic manner, characters to sentences and other complex expressions? It starts by assigning linguistic meaning (character) to semantically simple expressions of the language<sup>62</sup>. The next step is to go from the linguistic meaning (character) of simple expressions to the linguistic meaning complexes (including sentences). For this the theory specifies a general rule that derives the linguistic meaning of complex expressions from the linguistic meaning of their simple constituents and the syntactic rules of the language. The rule *can* have the following form<sup>63</sup>: if  $\alpha(e_i, e_j)$  is a complex expression and  $e_i$  and  $e_j$  are its immediate constituents and the meaning of  $e_i$  is a function whose domain contains the meaning of  $e_j$  then the meaning of  $\alpha(e_i, e_j)$  is the value of the meaning of  $e_i$  for the meaning of  $e_j$  as an argument:  $\mathbf{I}^*(\alpha(e_i, e_j)) = \mathbf{I}^*(e_i) (\mathbf{I}^*(e_j))$ .

This rule specifies the meaning function  $\mathbf{I}^*$  in a way that allows to derive the linguistic meaning (character) of any complex expression once we match up the variables in the rule with the constituents that correspond to them in each particular expression. Obviously if a semantics  $\mathbf{I}^*$  assigns meaning to every complex expression according to this rule, then it satisfies compositionality: the meaning of complex expressions *is* determined by the meanings of its constituents and the way they are combined since the meaning of the complex is the result of applying the meaning of one constituent (the functor expression) to

---

62 This includes all syntactically simple expressions but also some expressions that are syntactically complex but semantically simple e.g. idioms.

63 The rule given here is only one of the many types of compositional rules available to semantic theories.

the meaning of the other expression (the argument expression).

Using this rule and the assignment of meaning to simple expressions the semantics can derive the linguistic meaning (character) of any sentence. Once finitely many ways to combine expressions of the language are specified, and once each of the finitely many simple expressions of the language is assigned linguistic meaning if the semantics combines meanings according to this rule, then there is a systematic way to derive the linguistic meaning of complex expressions from the linguistic meaning of their constituents and the way the constituents are combined. Compositionality ensures that there is a systematic way to derive the linguistic meaning of complex expressions from the meaning of their constituents and their syntactic structure. This does not, yet, deliver truth-conditions, but standard semantic theories assume, and the anti-semantics skeptics deny, that the linguistic meaning of a sentence returns, for any given context, the truth-conditions of that sentence at that context. And this is what it means for a semantics to be systematic: that we can specify a rule (or a finite set of rules) by which the theorist can derive the truth-conditions of any sentence of the language.

## **2. Derivation of Truth-Conditions and Weak Compositionality**

As I said in the previous section the notion of weak compositionality was elaborated in the search for an account that accepts both the traditional purpose of formal semantics (namely that of modeling linguistic competence with natural languages by deriving truth value predictions for their sentences) and the conclusion put forward by radical contextualists (that for virtually any natural language sentence, its truth conditions are underdetermined by its linguistic meaning). The challenge to reconcile the two is a tremendous one, for obviously, if underdetermination is true then there is no systematic way for the semanticist to determine what sentences mean at a context, and thus no systematic way to predict the truth values of their respective utterances.

Weak compositionality promises to cushion the blow brought about by underdetermination for it promises to give the tools to build a systematic theory that allows



for pragmatic intrusion in the determination of truth conditions of natural language sentences. The idea is that by placing the requirement that semantic theories satisfy the weak version of compositionality for content, we can make room for pragmatic intrusion in a semantic theory (or rather a theory that aims to describe the fundamental meaning properties of a natural language by pairing sentences with their truth-conditions)<sup>64</sup>. In a properly weakly compositional semantics<sup>65</sup> it does not hold true that the content of a sentence depends on context only insofar as the contents of its constituents do. In other words, given a weakly compositional semantics of a language L if a sentence of L expresses context-dependent content its context-dependence need not be traceable to the content of its parts. Context determines the content of a sentence over and above determining the content of its constituents. In other words the content of a sentence can be determined at different contexts *in different ways* out of the content of its constituents (at the respective contexts).

Two things follow from these claims: (a) the contribution of the context to the content of a sentence cannot be fully specified or made explicit in the linguistic meaning (character) of its constituents and (b) in order to determine the content of sentences in a weakly compositional semantics it is not sufficient lay down the syntactic rules, assign contents to simple expressions and lay down the simple composition rules of the form like the one given above. Something else needs to be specified. The extra thing that must be specified is how context determines the content of sentences over and above its determination of the content of their constituents. I'll argue that there are good reasons to doubt that that the extra-contextual contribution can be specified in a systematic manner. In other words, I doubt that there are some general rules or procedures such that using those rules the theorist can derive for any sentence-context pair the extra contextual contribution to the content of that sentence at that context.

A discussion of arguments in favor of underdetermination and pragmatic intrusion will shed on why this is so.

---

64 Maybe it is a bit improper to use at this point the label "semantic theory" for this is an account that integrates pragmatics in the determination of truth-conditions.

65 By "a properly weakly compositional semantics" I mean a semantics that fails to be strongly compositional but is weakly compositional. From now, for reasons of economy, whenever I use "weakly compositional semantics" I refer to a properly weakly compositional semantics, unless I specify otherwise.

### 3. Open Texture and Systematicity

As I said above, a general strategy widely followed in semantic theorizing is the following: if a sentence *S* varies its truth-conditions across contexts of utterance, specify an element *x* that is the source of the variation and incorporate it in the semantics of *S*. According to supporters of underdetermination this strategy cannot be followed. Open-texture arguments purport to show that sentences vary their truth conditions along an open-ended number of un-systematic features, and that a theorist cannot specify all them. So, there is no systematic way for the theorist to determine in advance what truth conditions a sentence has relative to any arbitrary context.

According to some authors a sentence *S* can vary its content not only along one contextual parameter *x*, but along a multitude of features, in fact along an open-ended number of features. A theorist can specify a number of them that will determine the content of *S* relative to *some* contexts, but there will always be more, in fact an open-ended number of them, that are relevant for the truth-conditions of *S*. Virtually any natural language sentence is essentially open-ended in the sense that no set of specifiable contextual features can determine its content relative to every possible context of utterance. And thus there is no systematic way for the theorist to specify in advance for any possible context what *S* means at that context. In other words, contextual variation is un-systematic.<sup>66</sup>

Travis' arguments concerning (1) discussed in the previous chapter are of this type.

(1) "The leaves are green"

They rest on the intuition that there is a change in the truth-value of (1) without a change in the brute state of the leaves, since the leaves do not change color between the two utterances of (1). The argument aims to show that although the brute state of the leaves remains unchanged, the truth-conditions of (1) change if the interests and intentions of the

---

<sup>66</sup>Open texture arguments have a long tradition that goes back to Ziff 1972 and Searle 1978. See also Cohen 1985, 1986, Travis 1996.

conversational participants change<sup>67</sup>. In fact, one way to look at Travis' arguments is the following. The interests and intentions of the conversational participants make salient features that are relevant in the determination of the truth-value of (1) relative to that context. What makes (1) true at one context or another is not just the brute state of the leaves, but also whether they are green in a way that matches, or satisfies, the interests and intentions of the conversational participants. Under what conditions the leaves assume the color green is relevant for determining the truth value of (1) relative to a context. But the linguistic meaning of (1) is silent about the manner in which the leaves must be green.

This is also the gist of Searle's argument for semantic underdetermination. According to Searle, whether natural language sentences are true or false depends on an open-ended number of features. Some of these features are salient and obvious while others are not so. But when the semanticist specifies the truth conditions of a natural language sentence the specification obviously cannot be open-ended. But, argues Searle, any variation in features left out of this specification can bring about a change in the truth-conditions of the sentence.

To hedge his point he imagines several possible contexts of utterance across which (3) has not only different truth values but, allegedly, also different truth-conditions.

(3) "The cat is on the mat"

He starts by pointing out that when semanticists give the truth conditions of (3), they normally never explicitly mention an up-and-down orientation nor the presence of a gravitational field. But now, suppose that the cat and the mat are floating in outer space, or in a space lab where there is no gravitational field, but such that there is a contact point between them. Would we say that an utterance of (3) in this situation is true? The question about the truth-value of (3) in this context might have no clear answer, because it is not at

---

<sup>67</sup> Travis puts forward a plethora of examples similar to (1). Here is another one: Consider the sentence "The ball is round", and two cases of its use. *Case A*: What shape do squash balls assume on rebound? Pia hits a decent stroke; Jones watches. "The ball is round", she says at the crucial moment. Wrong. It has deformed into an ovoid. Jones did not say the ball to be as it was, so spoke falsely. *Case B*: Fiona has never seen squash played. From her present vantage point the ball seems a constant blur. "What shape is that ball?", she asks. "The ball is round", Alf replies; truly, since that it is the sort of ball a squash ball (and this one) is. It is not, e.g., like a very small rugby ball" (Travis 1996: 454). That is to say that a simple change in the interests and intentions of the conversational participants can lead to a change the truth-conditions of the sentence even though the state of the ball remains unchanged.

all clear in what relation the cat and the mat have to be, in the absence of a gravitational field, and thus in the absence of an intrinsic up-and-down orientation, in order for (3) to be true. So we might not know what to say about the truth value of (3). Then, it looks that the presence of a gravitational field is an essential feature of the truth conditions of (3): we cannot evaluate (3) as being true or false in the absence of a gravitational field. If (3) has a truth value only near the surface of the earth or in some other gravitational field, a theorist might be inclined to specify this as part of its linguistic meaning.

But this can't be so, argues Searle. It can't be part of the linguistic meaning of (3) that the cat is on the mat relative to a contextual field. For we can imagine further contexts at which (3) is true even though there is no gravitational field. He asks us to imagine that "we are strapped in the seats of our spaceship in outer space and we see a series of cat-mat pairs floating past our window. Oddly, they come in only two attitudes." From our point of view, and taking our bodies as giving an up-and-down orientation (eg. head-up, feet-down), either the cat is above the mat, or the other way around. "Which way is it now?", someone asks. "The cat is on the mat", comes the answer." Intuitively this time (3) can be true, even in the absence of a gravitational field.

Now the semanticist might be inclined to make part of the linguistic meaning of (3) reference to an up-and-down orientation. But according to Searle, even if the theorist makes explicit the reference to an up-down orientation in the semantics of (3) there are still an indefinite number of other features that she has to take into account. To support this claim Searle imagines yet more possible contexts of utterance for (3) where, intuitively, its truth-value of depends on further contextual features, still not made explicit in its semantics. For example, even if the reference to an up-and-down orientation is made explicit in the semantics of (3), its truth-value can still vary across contexts of utterance depending on the intentions of the conversational participants - or on features made salient by their intentions. "Suppose the cat and the mat are in the spatial relations [mentioned], at the surface of the earth, but that each, cat and mat, are suspended on an intricate series of invisible wires so that the cat, though slightly in contact with the mat, exerts no pressure on it.[...]. Suppose that the cat and the mat are part of a stage set. The wires are there to facilitate rapid movement of the props, as the cat has to be moved from chair to mat to

table. "Where is it now?" the director shouts from backstage; "The cat is on the mat" shouts his assistant." Intuitively, (3) is true. Presumably if the cat and the mat are in the spatial-relation just described but they aren't part of a stage set intuitively we would say that (3) is false in that context. So, it looks that whether the cat and the mat are part of a stage set is relevant for determining the truth value of (3) relative to a context of utterance. All this is taken to suggest that there is virtually no limit to what can be relevant for the truth or falsity of (3).

The moral of the story is that there is no end in sight to the list of features relevant for determining the truth value of possible utterances of (3). Given that there are an open ended number of possible contexts of utterance for (3), if speakers' intentions, interests and their world knowledge play a role in determining its truth-conditions, then there is no way for the theorist to establish beforehand for every context the truth conditions of (3). And if a theorist cannot determine under what conditions (3) is true relative to a context, a fortiori it cannot decide whether (3) is true or false at that context. This undermines the ability of semantic theories to pair sentences with their truth-conditions.

That is to say that if open texture arguments are right, then whether we judge (3) to be true or false depends on a highly complex and shifting space of intentions and interests. In what relation the cat and the mat have to be in order for (3) to be true depends on a wide web of interests, intentions and actions in ways that lie beyond the reach of formal theories. The point hedged by open texture arguments is that the semantic clause that gives the linguistic meaning of (3) is, or contains, a potentially open-ended disjunction. For example such a clause will tell us that (3) is true relative to a context iff the contextually salient cat is on the contextually salient mat in a gravitational field, *or* the contextually salient cat is on the contextually salient mat relative to a direction of orientation, *or* the contextually salient cat appears to be on the contextually salient mat, *and so on*. In other words, any semantic clause that states under what conditions (3) is true, will have to be infinitely long for it would have to state each of the possibly open-ended number of facts that are relevant for the truth-evaluation of (3). Then specifying the linguistic meaning (character) of (3) is an impossible task to fulfill. In fact, or so the skeptics claim, for virtually any declarative natural language sentence the clause that gives its linguistic meaning will be a potentially

open ended disjunction that states all the possible facts that are relevant for the truth evaluation of that sentence. That is, if formal semantic theories are in the business of providing specifications of truth conditions, any such specification will have to be an infinite disjunction. Which means that formal semantic theories are impossible to formulate.

The diagnostic that skeptics put for the failure of the project of formal semantics is that formal semantics rests on the unwarranted assumption that any contextual contribution to the truth conditions of a sentence “can be realized in the semantic structure of the sentence” (Searle 1987: 210) and that we can model natural languages with the help of formal ones. But if the arguments for underdetermination are true, then one has to accept that there is no systematic way to specify in advance how context contributes to the interpretation of complex expressions. Then, again, there is no systematic way to determine in advance what expressions mean in context, and a fortiori no way to test semantic theories by systematically generating truth value predictions about actual and possible utterances of sentences.

#### **4. Weak Compositionality and Systematicity**

This is where weak compositionality comes in. Some authors have tried to resist the radical conclusion drawn by skeptics and argued that systematic theories that describe the meaning properties of natural languages are still possible although not exactly as initially conceived.<sup>68</sup> We need theories that make room for pragmatics in the determination of truth conditions. But this brings another change. Theories that allow pragmatic intrusion through and through fail to be strongly compositional but they *can* be weakly compositional. Then, in order to accommodate pragmatic intrusion into truth-conditions, we need to give up the requirement that the theories about the meaning properties of natural

---

68 Recanati 2010, Pelletier and Pagin 2007, Pagin 2005, Pelletier 2003, Grandy 1990. Westerstahl 2010, Lasersohn 2012

languages satisfy strong compositionality in favor of the requirement that they satisfies weak compositionality. Pagin and Pelletier are explicit about how weak compositionality both makes room for pragmatic intrusion and saves systematicity. They claim that a weakly compositional semantics is “the middle ground [...] that integrates semantic composition and modulation [i.e. pragmatic intrusion]”, that is a “theory that both satisfies systematicity by obeying the intuitive content of the Principle of Compositionality and yet also allows for some general theory of language that [is] “between” pure literalism and radical contextualism (Pagin and Pelletier, 2007:32)”. Recanati concurs. It is worth quoting what he has to say about truth-conditional pragmatics (i.e. a weakly compositional account that accepts pragmatic intrusion) and systematicity in entirety:

“Here is the argument [that TCP is unsystematic]: in contrast to the contextual [determination of reference] to indexicals, modulation [i.e. pragmatic intrusion] is not driven by the linguistic meaning of words. Nothing in the linguistic meaning of words whose sense [content] is modulated tells us that modulation ought to take place. Modulation takes place purely as a matter of context, of 'pragmatics'; what drives it is the urge to make sense of what the speaker is saying. So modulation *is* unsystematic. If we allow it as a determinant of semantic content, we make it impossible to construct a systematic theory of semantic content. I grant the objector that modulation [pragmatic intrusion] is unsystematic. Still, I think it is easy to make room for it within a systematic semantics. In general, nothing prevents unsystematic factors from being handled systematically, by being assigned their proper place in the theory.”(Recanati 2010: 9) And he continues “Even though free pragmatic processes are allowed to enter into the determination of truth conditional content, still, in the non-minimalist framework I lay out [...] they comes into the picture as part of the compositional machinery. Semantic interpretation remains *grammar-driven* even if, in the course of semantic interpretation, pragmatics is appealed not only to [determine the reference] of indexicals but also to freely modulate the senses [contents] of the constituents in a top down manner” (Recanati 2010: 10)

I will argue that weak compositionality does not deliver the kind of systematicity

that we need. There are two senses in which pragmatic intrusion can be said to be systematic. One is that it can be handled systematically in the sense that it has a definite place and role in the theoretical machinery. The other is that pragmatic operations have a systematic contribution to the truth-conditions of natural language sentences: for any sentence-context pair the theorist can specify ahead what pragmatic operations are relevant for the truth-conditions of that sentence at that context. That is, pragmatic intrusion is systematic if there are a set of principles and rules that the theorist can use in order specify for any sentence at any context what pragmatics contributes to the truth conditions of that sentence at that context. This second way of understanding systematicity corresponds to our common way of understanding what it is for semantic theories to be systematic.

The paradigm of semantic systematicity is compositionality, we are often told. If a semantics is compositional, then it is systematic, which is to say that if a semantics is compositional then there is a systematic way for the theorist to work out the semantic values of complex expressions. For example, for a semantics that assigns characters as semantic values, satisfaction of character compositionality guarantees that there is a systematic way for the theorist to derive the semantic values of complex expressions, given syntactic analysis, assignment of semantic values to simple expressions and a specification of how semantic values (characters) combine (eg. by functional application or otherwise). Similarly for a semantics that assigns contents as semantic values, satisfaction of strong compositionality offers a theorist a systematic way to work out the content of complex expressions, given syntax, assignment of content to simple expressions and a specification of how semantic values (here contents) combine.

But weak compositionality does not deliver systematicity. Satisfaction of weak compositionality is not sufficient for there being a systematic way for a theorist to work out the content of complexes. In a weakly compositional semantics the theorist must also know how context determines the content of complexes over and above its determination of the content of simples. It is not enough that the theorist knows the content of simples (relative to a context) and how they combine, she must also know how context contributes to the content of complexes in a way that is not restricted to its contribution to the content of simples. That's because in a weakly compositional account context can determine the



content of a sentence over and above its determination of the content of constituents. In such an account in order to determine the content of a sentence relative to a context of utterance is not enough that constituents get assigned content relative to contexts, and that a syntactic analysis is assumed. The theory must also specify for any sentence and any context what is the extra contextual contribution of the context to the content of that sentence at that context. Unless the theory does this, it cannot determine the content of the sentence at that context. And obviously weak compositionality does not offer an account of that, for that is not what compositionality does. In other words, in order to save systematicity the appeal to the “trick [of] letting the composition function take the context as an extra argument” (Recanati: 2010, 46) won't be enough. For this doesn't solve the substantial problem: that of going in a number of steps and following some definite rules from the content of simple constituents and syntactic structure to the content of sentences. In order to be systematic a weakly compositional semantics has to specify a procedure, or a set of rules that determines for every sentence-context pair what the extra contextual contribution is to the interpretation of the sentence at that context.

Then, if the threat to the project of formal semantics is that truth conditions of natural language sentences vary in an un-systematic way, finding out that natural languages can be given a weakly compositional semantics is not going to help with systematicity. To guarantee systematicity a theory that accepts that context acts on the semantic content of sentences in a way that is not linguistically mandated must be supplemented with a systematic characterization of how context does that. A theory that accepts pragmatic intrusion must be supplemented with a systematic pragmatics.

Supporters of pragmatic intrusion on truth-conditions might say that that is precisely what we should do. A theory that delivers truth-conditions for sentences will consist of two parts. One part that aims to characterize systematically the meaning properties of linguistic expressions, and a second part that aims to characterize systematically how features that are not linguistically encoded bear on the truth-conditions of sentences.

The details and viability of a systematic account of how features not linguistically encoded in sentences play a role in their truth-evaluation are yet to be seen. But someone who accepts arguments to the effect that no standard semantic theory of natural language is

possible because truth conditions depend on highly complex and shifting web of intentions should, at least, be also skeptic towards the possibility of such an account. That is, someone who accepts the argument that truth-conditions are severely underdetermined by semantics should also doubt that it is possible to specify some general rules that will determine the right pragmatic operation for every sentence at every context.

Why? The short answer is that this is precisely what the arguments for underdetermination claim that theorists can't do. A theory that allows for pragmatic intrusion aims at pairing sentences with their truth-conditions (relative to a context). Given that the derivation of truth conditions is driven not only by what is linguistically encoded in sentences but also by the intentions and interests of the conversational partners<sup>69</sup>, then part and parcel of such a theory will be attribution of intentions, attitudes and practical inferences to the conversational partners. To derive truth conditions requires to find the right intentions: to ascribe to the conversational partners intentions in terms of which one can spell out the truth-conditions of a sentence at that context. For example, such a theory must be able to predict for any possible context of utterance under what conditions (3) is true or false at that context. It must be able to predict that (3) is true of some some green painted russet leaves, if (3) is uttered in a context where the conversational partners need the leaves for a photographic project, but that it is false if uttered in a context where the conversational partners need the leaves so to study their biology. The trouble with this move is that we're back where the anti formal semantics skeptics left us.

If the intentions and interests of conversational partners are relevant in the truth-evaluation of a sentence, then, in principle, "information from virtually anywhere and about virtually anything might have a bearing" on the truth-conditions of that sentence. Thus there is no non-arbitrary way for the theorist to exclude in advance any fact or any belief as being irrelevant for the attribution of some intention or another. This relates to a very well known problem when it comes to intentional explanations (i.e. explanations in terms of intentions and practical inferences attributed to agents). Because there is simply no limit on the amount of contextual evidence which might turn out to be relevant in ascribing one intention or another, intentional explanations are always defeasable. For example, even the

---

<sup>69</sup> Or as Recanati puts it, derivation of truth conditions is not driven solely by linguistic meaning but also by the need "to make sense of what the speaker is saying" (2010: 9)

most simple and plausible explanation of an action in terms of the intentions and practical inferences of the actor can be defeated if new contextual factors are deemed relevant to account for the actor's behaviour<sup>70</sup>. And any contextual factor may turn out to be relevant for the explanation of an action, which means that for any given action there might be a potentially open ended number of alternative explanations. Consider the following example that I borrow and adapt from Borg (2004a) of an intentional explanation: you see John putting money into a soda dispenser. A simple intentional explanation of this is that John *wants* to quench his thirst, and he *knows* that by inserting coins into the machine he will receive a drink. But this is defeasible. Imagine that you find out that Jack needs change. Then an alternative explanation of his action is that he wants some change and he believes that together with his drinks he will receive the change he needs. Or imagine that you see John's friend lying on the floor in distress as a result of a long and exhausting run. Then the best explanation for John's action is that he wants to get some drinks for his friend. If in order to match sentences with truth-conditions (relative to a context) a theory must attribute the right intentions and practical inferences to the conversational partners, then assignment of truth-conditions are defeasible in the same way in which intentional explanations are. But this is precisely the point made by skeptics like Searle and Travis: for virtually any natural language sentence there are an open-ended number of contextual factors that might be relevant for determining its truth conditions, and this endangers the capacity of formal theories to deliver truth-conditions.

To illustrate the point that weak compositionality is not sufficient to save systematicity I'll turn, now, to a widely used example by advocates of free pragmatic effects on truth-conditions.

We can easily imagine numerous contexts in which utterances of (4) express exactly what its conventional linguistic meaning says, namely that the ham sandwich stinks.

(4) The ham sandwich stinks.

Imagine, for example, that a group of people is sorting rotten food from good food, and one of them utters (4). Intuitively, (4) is true at that context iff the salient ham sandwich stinks. But with a little bit of imagination we can easily think of a different context in which an

---

<sup>70</sup> This is a point also pushed by Borg (2004a: 218-220)

utterance of (4) expresses the proposition that the person who ordered the ham sandwich stinks. For example, imagine that in order to maximize speed and efficiency restaurant workers tend to refer to their customers by the dish that they order. Then an utterance of (4) by one of them is true just if the person who ordered the ham sandwich stinks. The truth conditions of (4) at this context involve a person and not a ham sandwich<sup>71</sup>. In fact, we can think of many more such scenarios. Imagine, for example that health and sanitation workers customary refer to the food parlors in their jurisdiction by their best known dish. At the end of an inspection day, one of them utters (4). Intuitively, the truth-conditions of (4) at this context involve a food parlor rather than a ham sandwich or a person.

The first-blush reaction, and I believe the best reaction, that a defender of standard semantic theories can have in the face of this example, is to reject the semantic significance of the data: to claim that the utterance of (4) in the restaurant context involves deferred reference, and therefore it is *not a literal use* of (4). Since utterances that involve deferred reference are not literal uses of sentences, they need not be handled by the semantic machinery. They are on a par with other non-literal uses, like metaphorical uses or ironical uses, that no theorist expects to be treated semantically. Those who believe that deferred reference should be treated by semantics like to point out that the deferred referent enters into semantic relations like anaphora and ellipsis<sup>72</sup>. But as Stanley (2007: 206-207) shows, this is also the case with metaphor and irony, so this offers no good reason to believe that deferred reference is a semantic phenomenon. For example, in the case of metaphoric use of (5) the anaphoric “his” gets its value not from the literal content of “the pig in the next room” but from its metaphorical content.

(5) The pig in the next room wants his check immediately

Likewise, in (6) the metaphorical content is what is carried over in the ellipsis. Similar examples can be brought up for ironical uses.

(6) John is a pig, and Bill is too

---

71 This example is due to Recanati (2010: 167). A similar example involving “The ham sandwich left without paying” originates with Nurnberg (1995) and is used quite extensively in the literature that motivates weak compositionality. See for example Jackendoff 1997, Pagin and Pelletier 2007, Recanati 2004, 2010.

72 e.g. Jackendoff (1992)

Yet, no theorist would argue that metaphor and irony are semantic phenomenon, just because metaphoric content and ironic content can enter into anaphoric relations and can be carried over in ellipsis.

Advocates of pragmatic intrusion claim that the differences in truth-conditions of (4) are due to the fact that at different contexts the noun-phrase “the ham sandwich” makes different contributions to its the truth-conditions. This claim already invites giving up strong compositionality. The prospects of accommodating it within a strongly compositional semantics are slim. Lets accept that the semantic content of the noun-phrase (its contribution to truth-conditions) varies across contexts of utterance and its syntactic structure is:  $[[\text{the}]_{\text{DET}}[[\text{ham}]_{\text{N}}[\text{sandwich}]_{\text{N}}]_{\text{NP}}]_{\text{NP}}$ . There are two ways in which it can be accommodated within a strongly compositional semantics and neither of them seems promising. The first is to claim that the content of some of its constituents varies across contexts. It is implausible that the content of the determiner “the” varies, so maybe we should think that it is the content of the compound noun “ham-sandwich” that varies across contexts. This being a complex expression, here too we should look for the source of variation in the content of its constituents: “ham” and “sandwich” or some variable in its logical form. But it is highly implausible that either “ham” or “sandwich” have different contents at different contexts, such that in the restaurant context their content provides reference to the person who ordered the ham sandwich. Rather it is the entire NP “the ham sandwich”, or at least the compound noun “ham sandwich” that refer to the person who ordered the ham sandwich. (The first option seems more plausible, but we can also think of situations where the compound-noun is used with deferred reference, like in “Ham sandwiches are stingy”, used to refer to people who buy ham sandwiches.). Otherwise one can say that the syntactic structure of “the ham sandwich” is more complex than meets the eye. There are elements in its syntax which are not realized in its surface form and which are responsible for its contextual variation. A popular approach claims that compound nouns contain covert variables in their syntax and that context provides a value for these variables. This allows that a variety of relations to be expressed by a compound noun which need not traceable to the content of the constituents nouns. In other words, there is a free variable in the syntactic structure of compound nouns that stands for the relation that holds between the nouns. But it doesn’t seem that “the ham sandwich” varies its content across

the two contexts because the value of the relation that holds between the two nouns is different at the two contexts. Its contextual variation is of a different nature than, for example, the contextual variation of “child murderer” which in one context can be used to refer to someone who murders children, while in a different context to a murderer who is a child. Whatever hidden variables might be in the syntax of the noun-phrase they seem not to do this job. Finally, as a third way to save strong compositionality, and accept that the noun-phrase has varying contents across contexts of utterance, someone might claim that “the ham sandwich” is a lexicalized phrase (for example like “best man”) and thus although syntactically complex is, in some sense, semantically simple. But I believe that this misses the point as well. The use of “the ham sandwich” in the restaurant context does not seem to be like the use of “best man” in the context of a wedding. Lexicalized phrases like “best man” have a fixed meaning and express a constant content, while the content of “the ham sandwich” can vary freely with the context of utterance. Moreover, this can't be an acceptable solution for the very simple reason that many noun phrases that cannot possibly be lexicalized can, nevertheless, be used deferentially: “The ham sandwich in the corner wants more coffee”, or “The ham sandwich with tomatoes wants the bill”, or “The fried eggs punched the foreign accent”. So it looks that the best option a defender of strong compositionality has is to deny that deferred uses of (4) are to be treated semantically.

Obviously, advocates of free pragmatic effects on truth-conditions believe that the fact that (4) has different truth-conditions at the two contexts of use must be explained by theories that pair sentences with their truth-conditions. The account put forward by authors like Recanati (2004) and Pagin and Pelletier (2007) aims to do just that: (a) it takes deferred reference to be a phenomenon that must be accounted for by theories that pair sentences with their truth-conditions (relative to contexts of utterance) (b) claims that there are contextual ingredients in the truth-conditions of (4) that are provided through free (i.e. not linguistically mandated) pragmatic operations (c) aims to deliver the truth-conditions of (4) at both contexts (restaurant-context, sorting-out context) in a systematic manner.

Here is, briefly and in plain words, how Pagin and Pelletier account for the difference in truth-conditions of (4). According to them the linguistic meaning of each simple constituent expression of (4) determines together with the context of utterance the

content of that expression at that context. The result of combining, according to the syntactic structure of (4), the content of the simples, as determined by their linguistic meaning plus context, is the proposition that the ham sandwich stinks. This, according to Pagin and Pelletier, does not correspond to the truth-conditions of (4) at the restaurant context, since they do not correspond to the intuition that (4) is true in that context if and only if the person who ordered the ham sandwich stinks. According to them the truth conditions of (4) at the restaurant context are delivered by free pragmatic operations that takes the proposition that the ham sandwich stinks and yield the contextually appropriate proposition that the person who ordered the ham sandwich stinks. Loosely speaking we could say that there is a context-specific pragmatic operation that maps dishes into their orderers that is relevant for the truth-conditions of (4) at this particular context. So, what results from combining the literal content of the parts (i.e the content determined by their linguistic meaning and context), according to the syntactic structure of (4) is only an intermediate stop in the overall process of determining truth-conditions. (Within this account we can treat literalness as a limiting case: the pragmatic operation that delivers the truth-conditions of (4) in the context of selecting food, is (or can be modeled) as the identity function (see for example Recanati 2010:45)). So, according to this account the truth-conditions of (4) are determined relative to a context not only by its linguistic meaning and the (kaplanian) context of utterance, but also by free pragmatic operations. These pragmatic operations are context-specific in the sense that at different contexts, different operations can be at work; they are (or can be modeled as) functions that map propositions into propositions, and they are constrained solely by particular features of the context of utterance and whatever general facts coerce practical reasoning. Since this account allows context to freely provide material in the truth-conditions of (4), obviously it fails to satisfy strong compositionality. But, such an account, of a fragment of English, *can* satisfy weak compositionality.

The fundamental problem that I see with this account is that it eludes answering the objection from systematicity, namely that if pragmatic operations intrude into truth-conditions, derivations of truth-conditions will be un-systematic. Formally the idea of weak compositionality is fine, but the appeal to weak compositionality doesn't address the substantive problem regarding systematicity.

If the story of pragmatic intrusion is true, then sentence (4) can have an indefinite number of truth-conditions, each of them being the result of some pragmatic operation. If relative to a context the truth-conditions of (4) involve a ham sandwich, while relative to another context they involve a person, and relative to yet another context they involve a food parlor, it seems that the pragmatic operations that affect the truth-conditions of (4) are quite loose. What is definitely true is that the pragmatic operations through which deferred reference works are not constrained by the linguistic meaning of the expression. The entity that is the reference of an expression when used deferentially need not be something of which the linguistic meaning of an expression is true. We cannot truly predicate of a person the property of being a ham sandwich, nor can we truly predicate it about a food parlor<sup>73</sup>.

Again, the problem that advocates of pragmatic intrusion face, and that is not addressed by the simple appeal to weak compositionality is the following: if pragmatic operations affect truth-conditions, then in order to match sentences with their truth-conditions, the theorist must be able to specify ahead for each sentence at each context what pragmatic operation affects the truth-conditions of that sentence at that context. For example, the theorist must be able to specify, following certain rules and procedures, that relative to the restaurant-context, the relevant pragmatic operation for (4) is the one that maps dishes into their orderers, and not the one that maps dishes into food parlors. In other words, there must be a method for the theorist to predict, for any sentence and any context, what pragmatic operation applies to that sentence at that context, without making use of independent knowledge of the truth-conditions of that sentence at that context.

This task is more difficult than the task of merely restricting the class of pragmatic operations available. Restricting the class of pragmatic operations available for (4), that is, sorting out possible pragmatic operations from impossible ones is simply not enough for theories to pair sentences with their truth-conditions<sup>74</sup>. It is not enough to determine, for example, that there is no context at which (4) is true iff elephants are mammals, and therefore no pragmatic operation that delivers that as truth conditions for (4). For this is still different from determining in a principled manner exactly *what* pragmatic operation

---

<sup>73</sup> This point is also made by Stanley (2007)

<sup>74</sup> Some advocates of pragmatic intrusion seek to do exactly that. Pagin and Pelletier (2007:45) and Recanati (2010: 11, fn9)



determines the truth-conditions of *that* sentence at *that* context

Pagin and Pelletier (2007: 58-59) also suggest that selecting the right pragmatic process is similar with selecting the right referent of a demonstrative relative to a context. In both cases some form of pragmatic maneuvering is required. They point out that we don't have a general and fully satisfactory theory that will tell us how to predict what will be the most salient person, object or relation in a certain context. For example, even if we can specify the contextual parameters of a given context (i.e. who the speaker is, what the time, place and world of the utterance are) we still lack a theory that will predict the salience profile of the context: what will be most salient to the conversational partners at that given context of use. But even though the truth-conditions of sentences like (7) vary depending on the referent of the demonstrative "he", and determining the reference of "he" requires what is the most salient male in the context of utterance, most semanticists don't take this limitation to impinge on the systematicity of semantic theories.

(7) He is hungry

Pagin and Pelletier suggest that we should adopt the same attitude towards pragmatic operations on truth-conditions: we might not have a theory that can predict which pragmatic operation applies to a sentence on a given context of use. Nevertheless this should not be taken to be more worrisome than the lack of a theory that can tell us how to predict what will be most salient at a given context of use.

I find this suggestion highly problematic, for several important reasons. First, a theory that accepts pragmatic intrusion onto truth-conditions must predict for any given context of utterance what pragmatic operation applies to (4) at that context. Without this there is no way in which it can predict the truth-conditions of (4) at that context. On the other hand, there is a robust sense in which a semantic theory can specify the truth-conditions of sentences like (7) relative to a context, even if it can't predict the salience profile of the context (i.e. which is the most salient male at that context). Semantic theories can specify some general rule, of the following form, that determines the referent of the demonstrative given a proper specification of the context of utterance: an utterance *u* of "he" refers to the male that is most salient for the speaker at the time and place of the utterance. And from this it can derive truth-conditions of the following form: If *u* is an

utterance of the sentence “He is hungry”, and if the word “he” uttered in *u* refers to *x*, then *u* is true iff *x* has the property of being hungry at the time of the utterance. Semantic theories can give this kind of rules because they can specify the reference of “he” in terms of certain contextual parameters: the man that is most salient for the *speaker* at the *time* and *place* of utterance, is the referent of an utterance of “he”. What individual *is* the most salient one is out of the reach of semantic theories. There is no reason to suppose that semantic theories should tell us what determines that a particular utterance of “he” refers to whichever male person it does, other than that is the most salient male on that occasion of speech. In other words, there is no reason to suppose that semantic theories should tell us what particular individual satisfies the property of being the most salient male at a given context, nor there is any reason to believe that they should determine how conversational partners figure out who is the most salient male at a context of utterance. It is enough for semantic theories to specify this type general rules. It is an interesting empirical question how speakers and hearers determine the referent of demonstratives at a given context, but semantic theories are not in the business of modeling this. Semantic theories model those aspects of meaning that explain linguistic competence. But to identify the relevant features of a context, both to identify the contextual parameters (i.e. who the speaker is, what the time, place, world of the utterance are) and the salience profile of the context (i.e. what is most salient) requires more than linguistic competence. Facts pertaining to knowing who the speaker is, or which is the most salient male on a given occasion of speech, and so on, are not part of the linguistic competence with a language. The best that semantic theories can do (and should do) is to tell us how to determine the referent of the demonstrative, given an appropriate specification of the context of use.

But when it comes to pragmatic operations relevant for the truth-conditions of (4) we can't give the kind of general rules that we can give for (7). That is, we can't specify any general rule that will select the right pragmatic operation relative to a context, simply because pragmatic operations are not controlled by contextual parameters. There is no particular contextual parameter in terms of which the reference of “the ham sandwich” varies across contexts of utterance. Rather, its reference is determined by means of general pragmatic skills, and depends on global interpretational considerations.

If intentions, beliefs, interests and practical inferences are relevant for determining truth-conditions then nothing short of a complete theory of human rationality (or of practical reasoning) will be needed to deliver the truth-conditions of (4) relative to a context of utterance. And any semantics (i.e. any theory that pairs sentences with their truth-conditions) will, in fact, have to contain such a theory. I believe that irrespectively of whether one feels optimistic or not about the prospects of complete theory of human rationality this is an extremely heavy burden to place on semantic theories. In fact, one might wonder whether this leaves us in better position than that of the skeptics. It might turn out that the difference between skeptics and truth-conditional pragmatists is just one of inclination. Both believe that pragmatics intrude on truth-conditions, but the skeptics believe that pragmatics cannot be handled systematically, while truth-conditional pragmatists believe that it can <sup>75</sup>. Skeptics believe that only a theory that incorporates a complete theory of human rationality can assign truth conditions, but they are pessimistic about its prospects, so they believe that a theory that assigns truth-conditions to natural language sentences is a doomed to fail. On the other hand truth conditional pragmatists seem to be more optimistic about the prospects of a theory that will look into speakers' intentions, interests, beliefs and practical inferences and is able to tell a sentence's truth-conditions at any given context of utterance, as a function of these intentions, interests and beliefs. But in the absence of a detailed and positive account that can predict for any sentence and any context which pragmatic operation applies to that sentence at that context this optimism is just that.

To repeat, the problem is that advocates of pragmatic intrusion claim that for any sentence *S* and context *C* there is a pragmatic operation (a modulation function) that determines the truth conditions of *S* at *C*, but offer no general procedure, or set of rules, that can determine ahead of time the pragmatic operation (modulation function) relevant for *S* at *C*. If there are as many different pragmatic operations as there are sentence-context pairs, then nothing has been gained by simply stipulating that each sentence-context pairs is

---

<sup>75</sup> Interestingly, many authors who believe that it is possible to give a systematic and empirically adequate semantics for natural languages believe, on the other hand, that a systematic theory of communication is not possible (e.g. Fodor 1983, Cappelen and Lepore 2005). Their reason for this claim is identical with the one offered by skeptics for rejecting the possibility of formal semantics. They believe that any theory of communication requires a full theory of how interests, intentions, beliefs and practical inferences interact, and they don't believe that any such systematic theory is possible.

coupled with such a function. Rather what the supporter of pragmatic intrusion needs is a rule-based procedure for determining for any sentence-context pair the function that determines the truth-condition of that sentence at that context. But what is contextually relevant for the truth-conditions of (virtually) any given sentence seems to shift from context to context and does not seem to be something that can be given ahead of time. This is, in fact, the skeptics' main point about semantic underdetermination, and it is a point accepted by supporters of pragmatic intrusion. Then, postulating truth-conditionally relevant pragmatic operations, in the absence of some general rules that determine how they work, fails to explain how sentences have the truth-conditions that they have and how knowing these truth-conditions enables speakers to be linguistically competent. That is, in the absence of any general procedure that will determine for any sentence-context pair the pragmatic operation relevant for that sentence at that context appeal to such pragmatic operations fails to advance the explanatory enterprise. And, I argued, advocates of pragmatic intrusion should be skeptic about the availability of such general procedure, for the same reasons why they accept semantic underdetermination.

There is another way to put this. Assume, for the moment, that truth-conditional pragmatists are right in that speakers' intentions together with linguistic meaning and context determine the truth-conditions of sentences at contexts. Then it is trivially true that there is a function from sentence-context pairs to their corresponding truth-conditions. But claiming that there is such a function (as truth-conditional pragmatists do, when they claim that for each sentence-context pair there is a corresponding modulation function) is just restating what stands in need of explanation, namely how sentences have the truth-conditions that they have. To carry any explanatory weight truth-conditional pragmatics must (a) give a procedure by which it is possible to determine for each sentence-context pair the relevant pragmatic operation (modulation function) and (b) give an account of how particular pragmatic operations work. None of these challenges have been, yet, met.

It is fair to say that when it comes to systematicity weak compositionality offers no help. In the absence of a positive account that will be able to predict specific pragmatic operations, the systematicity challenge is not met. As yet, no one has put forward, as far as I know, a set of rules that can go over the intentions, interests and beliefs of conversational

partners and predict the truth-conditions of a given sentence at a given context. And there is one reason for the absence of such a general procedure: intentions and interests are formally intractable in that there is no way to predict what intentions, interests and beliefs are going to be relevant for the interpretation of a sentence at a context.

It is fair to conclude that weak compositionality fails to deliver some of its promised benefits. A weakly compositional theory that allows for pragmatic intrusion will fail to determine the truth-conditions in the absence of a general set of rules that can predict systematically the behavior of pragmatic operations. In other words, a theory that allows for pragmatic intrusion *can* satisfy weak compositionality and, nevertheless, fail to deliver a systematic derivation of truth-conditions. *A fortiori* such a theory will also fail to deliver testable predictions about truth-values of sentences in contexts. Therefore it will fail to offer an explanation of that core part of linguistic competence that is reflected in speakers' ability to tell how the world has to be in order for a sentence to be true.

## **5. Summing up**

In this chapter I considered the claim that weak compositionality can reconcile the kind of pervasive and radical context-sensitivity of natural language sentences, brought to light by skeptics, with some form systematicity. I argued that weak compositionality fails to deliver this.

It is claimed by some authors that natural languages exhibit pervasive and radical context-sensitivity: pervasive in the sense that it affects virtually any natural language sentence, and radical in the sense that it cannot be handled by fixing the values for a limited set of contextual parameters. Quite often it is argued that this claim threatens the project of giving a systematic semantics for natural languages. Lately though, many other authors have argued that there is no incompatibility between the acceptance of such pervasive and radical context-sensitivity and the possibility of systematic semantics, where by systematic semantics is understood a theory that systematically derives truth-condition for sentences at

contexts. Supposedly weak compositionality will help the theorist handle pervasive and radical context-sensitivity in a systematic manner.

Semantic theories are said to be systematic if they offer the theorist a method for determining beforehand the truth-conditions of sentences at contexts. In this chapter I argued that weak compositionality fails to deliver this kind of systematicity for theories that allow for pragmatic intrusion. As long as there is no method to determine beforehand what particular pragmatic function operates on each possible sentence-context pair, a theory will fail to systematically deliver truth-conditions even if it satisfies weak compositionality. And if we think of pragmatic intrusion as operations on meaning that “open possibilities for creative and unforeseen uses of expressions”<sup>76</sup> there is little hope for achieving that. I offered a few arguments for why we should be pessimistic about the possibility of a method that will predict for any sentence-context pair the relevant pragmatic function. The main contention is that such a method will require nothing less than a full theory of how interests, intentions, beliefs and practical inferences interact. Which means that a theory that derives truth-conditions will have to access, for any single derivation, a potentially open-ended space of intricate and highly specialized human interests and concerns. Finally, I pointed out that in the absence of a positive account that will determine for every sentence-context pair the relevant pragmatic function, the claim that there is such a function is just a restatement of what needs to be explained (namely how speakers' intentions and interests determine truth-conditions) and is not at all an explanation.

---

76 Bezuidenhout (2002:125)







## CHAPTER 5:

# Weak Compositionality and The Productive Character of Natural Languages

### 1. Introduction

It is widely held among semanticists that compositionality is an explanatory principle. Compositionality is expected to explain certain properties of natural languages and certain features of linguistic competence with natural languages<sup>77</sup>.

In fact, most motivations for wanting compositionality rely on such expectations. For example, compositionality is taken to explain how languages with infinitely many meaningful expressions can be learned by speakers with finite means, or how speakers competent with a language can understand new expressions that they've never encountered before. Also very often in the literature failure to satisfy compositionality is used as an argument against certain semantic theories. It is argued that such theories cannot be acceptable theories of a certain natural language (e.g. English) because they are not compositional. Presumably, because they fail to be compositional, they fail to explain certain interesting properties of that language and of linguistic competence with it<sup>78</sup>.

The plan of this chapter is the following. First I'll briefly present how compositionality explains the productive character of natural languages and of linguistic competence with them. I'll focus on the claim according to which if languages with finite vocabulary and a finite number of syntactic rules are compositional then they are learnable, and not on the stronger claim that only compositional languages are learnable. Then, I'll turn my attention to weak compositionality. I'll argue that because weak compositionality

---

<sup>77</sup> The claim that compositionality is an explanatory principle didn't pass unchallenged. For example Groenendijk and Stokhof (2005) and Dever (1999) take compositionality to be a methodological principle. But most semanticists expect compositionality to deliver some explanatory benefits with respect to certain features of linguistic competence.

<sup>78</sup> Probably the best illustration of this use of compositionality is found in the series of articles by Fodor and Lepore collected in their *The Compositionality Papers* (2002). There Fodor and Lepore reject a host of semantic theories on the grounds that they fail to satisfy compositionality, and thus fail to explain certain properties of natural languages languages.

allows that the manner in which meanings combine vary freely with the context of utterance it fails to explain the productive character of natural languages. I'll give several arguments of why this is so. Finally, I'll consider a possible fix: one according to which weak compositionality delivers its explanatory benefits only together with additional constraints on the form of semantic rules. There are two ways in which we can constrain the form of semantic rules: one will be of no use, but the other will alleviate the worries about the explanation of the productive character of natural languages. However, as I'll argue this fix undermines the motivation for wanting weak compositionality in the first place: that of making room for pragmatic intrusion in the determination of content. Moreover, as I'll discuss in the last section, the idea of weak compositionality is still open to another worry: it presupposes an implausible account of the semantic significance of syntactic rules.

## **2. Compositionality and Productivity Facts**

It is also widely held among natural language semanticists that compositionality explains the productive character of linguistic competence, that is, speakers' ability to produce and understand an open-ended number of complex expressions. More precisely compositionality will explain how it is possible that linguistic competence with natural languages is infinite in scope when otherwise speakers of natural languages have finite cognitive means: finite memory, finite time, finite computational capabilities, etc. Generally claims about the productive character of linguistic competence come under different guises. One is an observation about *learnability*: given that natural languages have more meaningful complex expressions than what can be learned one-by-one but speakers can, nevertheless, learn them, it means that natural languages have some properties that speakers exploit in order to learn the meaning of their complex expressions. Sometimes claims about productivity are put in terms of understanding: speakers of a given natural language can, in principle, understand any of the infinitely many meaningful complex expressions of that

language. The other guise that claims about productivity take is the observation about *novelty*: speakers of a given natural language commonly and constantly understand complex expressions that they have never encountered before. This stands in need of explanation: there must be a property that natural languages have and that speakers exploit such that speakers reliably understand complex expressions that they have never encountered before. Sometimes, claims about novelty are supplemented by the observation that linguistic communication with new sentences has a high rate of success. This shows that both speakers and hearers understand new sentences in that they pair them with the same meaning<sup>79</sup>. Learnability and novelty (in short facts about productivity) are observable features of linguistic competence and therefore part of the data that semantic theories aim to explain. That compositionality plays a role in their explanation offers a good reason for requiring that semantic theories satisfy compositionality.

### 3. The Learnability Argument

So how can we explain speakers' capacity to understand a potentially infinite number of meaningful sentences given their finite cognitive means? This is how the learnability argument in favor of compositionality, generally, gets started. Compositionality gets us out of an apparent conundrum by squaring together the following facts about natural languages and linguistic competence: (a) natural languages have a potentially infinite number of meaningful sentences<sup>80</sup> (b) natural languages are learnable, and their speakers have the capacity to understand any of the infinitely many meaningful sentences (c) human speakers of natural languages have limited cognitive capacities, and can learn one-by-one only a limited number of meaningful sentences

If speakers were to learn the meaning of sentences of a given natural language one

---

79 See Pagin (2003)

80 Some authors doubt that natural languages have an infinite number of sentences (e.g. Pullum and Scholz 2012). But notice that the claim that natural languages have an infinite number of meaningful sentences (which might be, after all, just a useful idealization) is not really needed to motivate compositionality. What is needed is, as Grandy (1990: 556-560) points out, that natural languages have more meaningful sentences than what can be learned one-by-one by human speakers. And by any plausible estimation they have.

by one, given that there are an infinite number of meaningful sentences, they would never accomplish the task of learning that language. What explains that they nevertheless learn a given language is that they learn a finite set of meaningful simple expressions and a finite set of rules, and that the semantics of that language is compositional. This offers the rudiments of an abstract explanation of learnability of natural languages.

There are, in fact, two ways to formulate the learnability argument for compositionality. One is to claim that compositionality is a *necessary* condition for learnability: no language with an infinite number of meaningful sentences can be learned unless its semantics is compositional<sup>81</sup>. Then there is no language with infinitely many meaningful complex expressions that is learnable but is not compositional. The other is to claim that compositionality is a *sufficient* condition for learnability: if a language has a compositional semantics, then it is learnable. As Szabó (1995) emphasizes the argument in terms of sufficiency is that if a language is compositional we can understand any of the infinitely many meaningful complex expressions by learning the meaning of its simple constituents and its syntactic structure. The argument in terms of necessity adds that this is the only way in which we can understand a language with an infinite number of meaningful complex expression.

The argument in terms of necessity is often taken to show that natural languages *are* indeed compositional. Since speakers cannot accomplish the task of learning them by learning the meaning of their sentences one at a time, it must be that there are finitely many features of these languages which are learnable and which speakers exploit in learning them. Natural languages have finite vocabularies and finite number of syntactic rules which speakers can learn and exploit. For speakers to be able to learn an entire language by learning a finite number of meaningful expressions and a finite number of syntactic rules their semantics *must* be compositional. If we understand compositionality as a necessary condition for learnability then the fact that natural languages are learnable, is taken to show that they must have a compositional semantics. Here is an often quoted

---

81 Except for languages in which all of the infinitely many sentences have the same meaning. In that case, if one learns that they mean the same, and one learns what they mean, one has accomplished the task of learning the language. If we are to be strict we should speak about languages with infinitely many meaningful sentences that can express infinitely many meanings. (Pagin and Westerstahl 2010:265)

passage from Davidson where compositionality is understood this way: “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. For suppose that a language lacks this feature [compositionality]; then no matter how many sentences a would-be speaker learns to produce and understand, there will remain others whose meanings are not given by the rules already mastered. It is natural to say such a language is unlearnable”. (Davidson: 1965, 8)

The argument is more often understood in terms of sufficiency. In these terms the argument claims that if speakers have the capacity to know the meaning of simple expressions of a language and the syntactic rules of the language, then suffices that the language has a compositional semantics to guarantee that speakers can learn the entire language. In other words, if speakers can learn the meaning of a finite set of simple expressions and a finite set of syntactic rules of a language and if the language is compositional then the language is learnable by speakers with finite means. But rather obviously, if we understand the learnability argument in terms of sufficiency, the fact that natural languages are learnable doesn't guarantee that they have a compositional semantics. Nevertheless compositionality may still be a plausible explanation for why natural languages are learnable by speakers with finite resources. Given that compositionality offers a plausible explanation for learnability of natural languages, this is a good enough reason for requiring that their semantics satisfy compositionality.

In what follows I'll focus on this weaker, but more common, understanding of the relation between the learnability (and other productive features) of natural languages and compositionality: if a language has a compositional semantics, and its vocabulary and syntax are learnable by speakers with finite means, then the entire language is learnable by speakers with finite means.

#### 4. The Explanatory Power of Compositionality

It should be fairly easy to show that compositionality is not sufficient for learnability. Compositionality establishes only a functional relation between, on one hand, the meaning of complex expressions and, on the other hand, the meaning of their constituents and their syntactic structure. A semantics satisfies compositionality if there is *some* function such that for every meaningful complex expression it takes the the meaning of its constituents and its syntactic structure and returns the meaning of the complex. But compositionality is silent with respect to what that function is. That is to say that compositionality does not eliminate *any* choice of function as the composition function employed by a semantics. Then it is fairly easy to think of languages that satisfy compositionality but are unlearnable. Any semantics that satisfies compositionality but employes an unlearnable (i.e. strictly speaking, an uncomputable) composition function is unlearnable. Then, the language defined by that semantics is unlearnable, although it is compositional. In other words, a semantics **I** satisfies compositionality if there is a function  $f$  such that the meaning of complex expressions are an  $f$ -function of the meaning of to their constituents and syntactic structure. But  $f$  need not be computable. If  $f$  is not computable then neither is **I**. Simply put the mere fact that a semantics **I** satisfies compositionality does not guarantee that it is learnable because it doesn't guarantee that its composition function  $f$  is computable. The meaning of its complexes might be a (non-computable) function of the meaning of their parts and their syntactic structure.

This brings forward an important problem concerning compositionality, namely whether the explanatory purposes of compositionality can be served by a mere functional relation. For if compositionality requires only that there is *some* function from the meaning of parts plus syntax to the meaning of complexes, such that *any* function will do, there are good reasons to doubt that compositionality can deliver any of its expected benefits. And, in fact, there are authors who doubt this and suggest that compositionality has explanatory power only “together with a fuller specification of what is required by the relation “is a function of” (Partee: 1995, 154). The suggestion is that in order for compositionality to have explanatory value with respect to those semantic properties that it seeks to explain, we

should ask that the composition function satisfies some further properties. That is, we should require that only certain functions are permissible composition functions and we should give a specification of that class of permissible composition functions. For example, following this line of thought, some authors suggest that we should require that the composition function(s) employed by a natural language semantics is, at least, computable (e.g. Grandy 1990: 560)). Then any semantics that satisfies compositionality and this requirement concerning the functional relation established by compositionality is learnable.

The requirement that the composition function is one that we can master is only a minimal requirement. Some authors have argued that even stronger constraints should be placed on what kind of functions are acceptable composition functions if we want compositionality to deliver its expected explanatory benefits<sup>82</sup>. And in fact, there are good reasons for placing such constraints, independently of considerations regarding the explanatory power of compositionality. When the theorist gives a semantics for a natural language she must specify how meanings combine and she does that by stating a set of rules that determine just that. And each rule specifies one or more ways in which meanings combine. (More about this bellow) Obviously there are an indefinite number of ways in which meanings can combine, but some of them cannot possibly be ways in which meanings of natural languages combine. So the theorist must require that when it comes to natural languages, the composition function (that is, the manner in which meanings combine) is selected from a class of acceptable functions. And it is part of her job to delineate the class of permissible or acceptable composition functions for natural languages.

Then the idea is that compositionality is explanatory useless unless we put some constraints on what kind of functions are acceptable composition functions. I'll illustrate this point with the help of a little argument due to Szabó (2000)<sup>83</sup>.

---

82 e.g. Szabó 2000

83 I borrow from Szabó the manner or reasoning, but I changed the example so to suit my purposes better.

## 5. English and Inverted English

Assume that English (or a fragment of English) is character compositional. Now, consider the following language, call it Inverted English, which has the same vocabulary and the same syntax as English, and its sentences have the same meaning as their English counterparts with the important exception of (3) and its synonyms and of (4) and its synonyms:

(3) John kissed Mary

(4) Mary kissed John.

Sentence (3) and sentences synonymous with it, express in Inverted English what sentence (4) expresses in English, and (4), and sentences synonymous with it, express in Inverted English what (3) expresses in English. More clearly, the Inverted English sentence (3) and its synonyms, express the proposition that Mary kissed John. And the Inverted English sentence (4), and its synonyms, express the proposition that John kissed Mary. The meanings of (3) and (4) (and their respective synonyms) in Inverted English are obtained from their meanings in English through a function that interchanges their meanings while it leaves the meaning of every other expression as it is in English. In other words, the meaning function of Inverted English can be represented as a composite of the meaning function of English and a function (defined on the values of the meaning function of English) which leaves every meaning as it is in English except for (3) and (4) and their respective synonyms for which it interchanges their meanings. Then the meaning function of English and Inverted English agree on all assignments except for (3) and (4), and their respective synonyms.

Given the initial assumption that English is compositional, it can be shown by a *reductio* that Inverted English is compositional too. Assume that Inverted English is not compositional. Then there are two sentences  $S1$  and  $S2$  of Inverted English that are not synonymous although they share the same syntactic structure and have their constituents are pairwise synonymous. From the definition of the meaning function of Inverted English it follows that  $S1$  and  $S2$  are synonymous in English just in case they are synonymous in



Inverted English. Then we can conclude that *S1* and *S2* are not synonymous in English, despite the fact that they share the same syntactic structure and their constituents are pairwise synonymous. From this, it follows that English fails to be compositional, which contradicts the initial assumption that English is compositional. So, Inverted English does not fail to be compositional. ■

At the first glance it might be surprising that both English and Inverted English are compositional. But at a closer look it is not surprising at all. The semantics of Inverted English is a function that agrees with the semantics of English for every expression, *except* for (3) and (4) and their respective synonymous. We can say that the semantics of English and Inverted English differ because they employ different composition functions. For example, we can represent the composition function of Inverted English as a composed function  $f_2 \circ f_1$ , where  $f_1$  is equivalent with the composition function of English, and  $f_2$  is a function that takes the values of  $f_1$ , and yields meanings. For most expressions,  $f_2$  yields the same value as the identity function except for (3) and (4) (and their respective synonyms), for which it interchanges their meaning: to (3) and its synonyms it assigns the meaning that (4) has in English and to (4) and its synonyms it assigns the meaning that (3) has in English.

This result should not come as a surprise. Again, compositionality doesn't state anything whatsoever about what functions can be employed as composition functions. It is perfectly compatible with compositionality that there are two languages  $L_1$  and  $L_2$  such that they have the same vocabulary and syntax, but their respective semantics' employ different composition functions,  $f_{L1}$  and  $f_{L2}$  respectively, and that each of them satisfies compositionality. It can hold true about the semantics of  $L_1$  that the meaning of complex expressions of  $L_1$  are an  $f_{L1}$  function of the meaning of their constituents and of their syntactic structure. And it can hold true of the semantics of  $L_2$  that the meaning of complex expressions of  $L_2$  are an  $f_{L2}$  function of the meaning of their constituents and their syntactic structure. Then precisely because their semantics' employ different composition functions there is at least one complex expression of  $L_1$  and one complex expression of  $L_2$  that differ in meaning although they have identical syntactic structure and pairwise constituents with equivalent meanings.

Returning to English and Inverted English, it should be obvious that Inverted English is learnable. Anyone who masters English will have already learned Inverted English by the time they finished reading the little argument given above. Nevertheless, Inverted English is odd, to say the least. And probably any theorist of natural languages will say that Inverted English cannot possibly be a natural language. Presumably they would say that Inverted English violates some important facts about how syntax and semantics of natural languages interact. Syntactic structure is not interpreted uniformly in Inverted English, and there is no correlation between the syntactic position given to constituents of sentences and their role in the semantic interpretation of the sentence. For all Inverted English sentences, with the syntactic structure  $[[e_1]_{NP} [e_2]_V [e_3]_N]_{VP}]_S$  the noun phrase that is immediately dominated by the top node is interpreted as the semantic subject of the sentence, and thus as the agent of the action described by the verb, *except* for (3) and (4) where it is interpreted as the semantic object, and thus as the patient of the action. And for all Inverted English sentences, with this syntactic structure the noun that is immediately dominated by the verb phrase is interpreted as the object of the sentence, *except* for (3) and (4) where it is interpreted as the subject of the sentence. In Inverted English it is not true that the steps of syntactic combinations and semantic combinations are the same<sup>84</sup>. So, a natural language semanticist might argue that there are good reasons to believe that the composition function by which meanings of Inverted English combine is not acceptable from the point of view of natural languages<sup>85</sup>. For example, some semanticists (Montague (1970), Partee (1979) and their followers) would add constraints that introduce a certain uniform mapping between syntactic categories and types of meaning, in the sense that (a) all expressions in the same syntactic category have the same kind of meaning or semantic category and (b) all syntactic rules that have the same input and output grammatical categories, are interpreted by the same compositional semantic rule. Operating with these

---

84 Someone might claim that it is not true that English and Inverted English have the same syntax. But that is just not accepting the stipulations regarding Inverted English. Also, someone might reject that Inverted English is compositional because intuitively in Inverted English the meaning of “loves John” is not a function of the meaning of its constituents. But by stipulation “loves John” means in Inverted English what it means in English, and if its meaning in English is a function of the meaning of its constituents, then so it is in Inverted English. Finally, both (3) and (4) of Inverted English are in active voice.

85 Others, (e.g. Szabó 2000:501) might argue that Inverted English cannot possibly be a natural language because it can't be learned as a first language. The only way to learn Inverted English is to learn English and then learn the rule that introduces the permutation on the sentence meanings.

constraints one can rule out the composition function of Inverted English, as an acceptable function for natural languages.

The job of the natural language semanticist, then, is to determine what ways of combining meanings are usable by semantic theories of natural languages. And given that ways of combining meanings are introduced by semantic rules, this is equivalent with saying that part of the job of the semanticist is to determine what compositional rules are acceptable rules for natural languages. Presumably, selective permutation of meanings is not an acceptable operation on meanings, nor is the dedicated semantic rule of Inverted English that introduces it a rule that can be found in natural languages. Deciding what compositional rules are acceptable rules for natural languages is, in fact, part and parcel of the job of the natural language theorist. The natural language theorist is interested, among other things, in giving a characterization universal properties of natural languages. The more narrow she characterizes the class of possible natural languages the more substantial and more informative her semantic theory is. To do this, she must decide what syntactic ways of combining expressions are possible in natural languages, what kind of meanings can be assigned to expressions (or rather what kind of entities can fruitfully be used as meanings), and what compositional rules are acceptable rules for natural languages. There are many considerations that can guide the natural language semanticist in restricting the class of permissible composition functions, just as there are many considerations that can help her in restricting the class of possible syntactic rules of natural languages.

The take home message is that unless further constraints are placed on the functional relation established by compositionality (i.e. constraints on what functions are allowed) compositionality is explanatory useless with respect to interesting properties of natural languages. Moreover, independent of worries concerning the explanatory benefits of compositionality, the natural language semanticist has other motivations for requiring such constraints, as they are nothing but constraints on the form of semantic rules.

## 6. Weak Compositionality: Promises and Disappointments

In this section I'll argue that weak compositionality fails to deliver some of its expected explanatory benefits, namely that it fails to explain the productive character of natural languages. More precisely, weak-compositionality allows for context-sensitive semantic rules that introduce composition functions which vary freely with the context of utterance. I'll argue that it is impossible to state in a finite manner such rules. To make things worse appeal to restrictions and constraints on the acceptable composition functions will be of no help. In short: the type of semantic rules employed by a weakly compositional semantics that allows the composition function to vary with the context are in conflict with the productive character of natural languages

In order to keep things clear, let me rehearse a basic distinction introduced in chapter one: that between *rules* and *operations*. This distinction applies both at the level of syntax and at that of semantics. *Syntactic rules* state how expressions of a language combine to form larger grammatical expressions. Here is an example of a syntactic rule: if  $e_1$  is an expression of category NP and  $e_2$  is an expression of category VP then concatenating  $e_1$  and  $e_2$ , in this order, results in an expression of category S. This rule introduces one specific type of *syntactic operations* by which expressions combine, namely concatenation, but other rules can introduce other types of operations. *Semantic rules* state how the meaning of expressions with a certain syntactic structure is obtained. For example, one commonly encountered type of semantic rule is the following: if an expression has the syntactic structure  $[[e_1]_{NP} [e_2]_{VP}]_S$  then the meaning of  $S$  is obtained by functionally applying the meaning of  $e_2$  to the meaning of  $e_1$ :  $\mathbf{I}(S) = \mathbf{I}(e_2)(\mathbf{I}(e_1))$ . Each semantic rule specifies one or more *semantic operations* by which meanings combine. This rule, for example, introduces one semantic operation: functional application, but other rules can specify other operations. In a nutshell: each syntactic rule specifies that expressions of certain syntactic categories can combine to form expressions of a certain syntactic category, and determine the syntactic operation by which they combine. And each semantic rule specifies an operation that determines the meaning of complex expressions given the meaning of their constituents and the syntactic rule that builds them. In a sense, semantic rules interpret syntactic rules.

My worry concerning weak compositionality concerns whether it is possible to formulate weakly compositional semantic rules that, on one hand, are learnable and, on the other hand, can accommodate the kind of thorough-going context-sensitivity that is not amenable to explanation in terms of a fixed set of contextual parameters. Semantic rules determine the manner in which the meaning of constituent expressions combine to give the meaning of complexes, and standardly each semantic rule specifies a unique manner of combining meanings. But weakly compositional semantic rules will specify for each syntactic structure that they interpret several manners of combining meanings, that is several semantic operations. Then a single syntactic structure will contribute in more than one way to the interpretation of complex expressions, and its contribution will vary with the context of utterance. Given that there are indefinitely many contexts of utterance, a weakly compositional semantic rule that interprets a single syntactic structure will, then, specify indefinitely many ways of combining meanings. If these ways of combining meanings vary across contexts not as a function of a fixed set of contextual parameters, then learning such rules is impossible.

Let me develop this point and also put this from a slightly different angle. The problem with a properly weakly compositional semantics is that it employs context-sensitive semantic rules. In a weakly compositional semantics it holds true that the context-sensitivity complex expressions need not be traceable to its constituents: it can be traced to the semantic rule that determines how the meanings of constituents combine. For example, in a weakly compositional semantics it holds true that a sentence like (1) can express different contents at different contexts, although the content of its constituents stay the same across contexts of utterance.

(1) John runs

Then, relative to a context the content of (1) is determined by more than just applying the content of the verb-phrase “runs” to the content of the noun-phrase “John”. In different contexts the content of (1) is a different function of the content of the noun-phrase and the verb-phrase. Which is to say that the way in which the content of “John” and “runs” combine so to give the content of (1) varies across contexts of utterance.

To see this better consider, again, the formulation of the principle of weak

compositionality:

For every syntactic rule  $\alpha$  there is a function  $f$  such that for any expressions  $e_i, e_j$  and for any context  $C$ , if  $\alpha(e_i, e_j)$  is meaningful then  $\mathbf{I}(\alpha(e_i, e_j), C) = f(\alpha, \mathbf{I}(e_i, C), \mathbf{I}(e_j, C), C)$

Notice that by schönfinkelizing (currying)  $f$  this is equivalent with:

For every syntactic rule  $\alpha$  there is a function  $f$  such that for any expressions  $e_i, e_j$  and for any context  $C$ , if  $\alpha(e_i, e_j)$  is meaningful then  $\mathbf{I}(\alpha(e_i, e_j), C) = f(C)(\alpha, \mathbf{I}(e_i, C), \mathbf{I}(e_j, C))$

In plain words this says that at different contexts of utterance, different functions  $f(C)$  combine the contents of constituents into the content of the complex. For any single syntactic structure *the way* in which the content of the constituents of expressions with that structure combine can vary from context to context. Then, given that there are an indefinite number of contexts of utterance it follows that for any single syntactic structure there can be an indefinite number of composition functions. Which is to say that in a weakly compositional semantics, for any single syntactic structure the semantic rule that interprets it will specify an indefinite number of ways in which contents of constituents can combine to give the content of complexes with that structure. But a semantics that allows for an indefinite number of semantic operations that vary freely with the context will, obviously, fail to explain the productive features of linguistic competence: learnability and novelty. If there are an indefinite number of semantic operations which vary freely with the context, then speakers will have to learn an indefinite number of them one at a time.

A different way to drive the point about explaining productivity facts is the following. Given that semantic rules specify the operations by which meanings combine, if semantic operations are potentially different in different contexts given that there are an indefinite number of contexts, semantic rules must specify a potentially indefinite number of semantic operations. More precisely semantic rules must specify a potentially infinite disjunction of semantic operations. Obviously such rules are not finitely storable, so they are unlearnable. Let me elaborate a bit. Consider sentences with the structure  $[[NP] [VP]]$ . Weak compositionality allows that an expression with this structure varies its content across contexts of utterance although neither the NP nor the VP vary their content across contexts. And context-sensitive rules allow that the manner in which the content of the NP and the content of the VP combine into the content of  $S$  varies freely with the context of

utterance: at different contexts different operations combine the contents of the NP and that of the VP. Since these operations can vary freely with the context (i.e. not as a function of a fix set of contextual parameters) and the rule must introduce each and all of them, then the form of the the semantic rule must be an infinitely long disjunction. Such rules cannot be stated in a finite manner, and therefore they are not learnable<sup>86</sup>.

An equally bad option is to give for any single syntactic structure an indefinite number of rules and have each rule specify a unique semantic operation (a unique way of combining meanings). A single syntactic structure will be, then, interpreted in different contexts by a different semantic rule. Then for the syntactic structure [[NP][VP]]S has associated at every context a different rule semantic rule that determines how the contents of the NP and the VP combine so to give the content of S. Rather obviously such a semantics fares no better as it is not finitely statable either.

Nor will limitations on the functional relation established by compositionality help. As I explained, compositionality will explain little if no limitations are placed on the functional relation that it establishes. That's why we are advised to demand something more than there is a function from the semantic values of simples plus syntactic structure to the semantic values of complexes. As I explained in the previous section, we should always require that the composition function(s) satisfy some additional properties. But notice that this advice won't help when it comes to weak compositionality. The fact that there are an indefinite number of composition functions that can vary freely with the context of utterance renders such constraints useless. For example, notice that even if we require that each composition function is computable if there are indefinitely many of them, and they vary freely with the context of utterance, productivity features will still remain not guaranteed. As a way of putting more flesh on this consider an example.

Theorists who accept context-sensitive rules are motivated in doing this by their desire to make room for pragmatics in the derivation of truth-conditions. They believe that

---

86 In fact a theory that employs this kind of rule runs into the same kind of trouble as the one raised by the skeptic. Then obviously, theories that employ context-sensitive semantic rules cannot offer any relief from skeptics' challenge. And getting this relief, as I emphasized in the previous chapter, is the main motivation for weak compositionality. Skeptics claim that because virtually any natural language sentence vary its truth-conditions along an open-ended number of contextual features, the clauses that state their truth-conditions would have to be infinitely long. Cohen (1985:134) formulates his skepticism towards formal semantics in this manner.

a sentence (including indexical-free sentences) can express an indefinite number of propositions. Of course, they agree there are restrictions on what propositions a sentence can express. Although the English sentence (1) can express an indefinite number of propositions, it can't express *any* proposition.

(1) John runs

For example it cannot express the proposition that kangaroos have tails - there is no context of utterance such that (1) is true at that context if and only if kangaroos have tails. In other words, not everything goes. The propositions in question all have to be, in some sense, compatible with what is linguistically encoded in (1), that is with the linguistic meaning of its constituents and its syntactic structure. This is to say that not every way of combining semantic values are acceptable in English. There is no permissible composition function that takes the content of "John" and of "runs" at a context, and syntactic structure of (1) and yields the content that kangaroos have tail. Obviously, even the most radical truth-conditional pragmatists accept that speakers don't have absolute liberty to tinker with the meanings of the sentences they use; that speakers are not the absolute masters of their own words. Then even truth-conditional pragmatists will point out that even if there is pervasive and thorough-going context sensitivity not every function is a permissible composition function. They would, then, propose to distinguish those functions which are permissible composition functions for English from those that are not<sup>87</sup>. But this won't help. For if they accept that a sentence can express an indefinite number of propositions, and this variation is not due to any variation in the content of its constituents, then they accept indefinitely many composition functions. For every context of utterance there is a different composition function that determines the content of (1) at that context. And if this variation can't be handled by fixing the values for a limited set of contextual parameters then the only way to introduce these functions is to list them. But this is impossible, since there are indefinitely many of them. In fact, I believe that what is needed to domesticate context sensitive rules is to specify a fix set of contextual parameters in terms of which the composition functions defined by rules vary. That is to say that although semantic operations vary with the context of utterance, they vary according to some fix set of

---

87 See Recanati (2004: 6) and (2010: 10-11) for this.



contextual parameters. I'll discuss this option in the next section.

Finally, advocates of weak compositionality might seek refuge in rules of the following form: if  $e_3$  is a complex expression and  $e_1$  and  $e_2$  are its immediate constituents, then the content of  $e_3$  at any given context  $C$ , is the result of combining the content of  $e_1$  at  $C$  with the content of  $e_2$  at  $C$  in the way relevant at  $C$ <sup>88</sup>. But this move won't help a bit. Such a rule is useless to the natural language semanticist because of its total unspecificity about how contents combine. More precisely, the rule is useless when it comes to derivations of truth-conditions. In order to derive truth conditions the theory must also specify for each context the relevant way in which contents combine at that context.

Before I discuss domestications of context-sensitive rules I must fend off a possible source of misunderstanding. The arguments from learnability, novelty, and other productive features of linguistic competence, are about the repeatable and context-independent aspects of meaning. In other words they are about linguistic meaning. If anything they offer support for the compositionality of linguistic meaning (what I have called the principle of character compositionality). On the other hand weak compositionality, and the context-sensitive rules it allows, are about content. So one might wonder whether there is any conflict between context-sensitive semantic rules and the productive features of natural languages. But there is. For if it is impossible to state in a finite manner the semantic rules of a properly weakly compositional semantics that assigns content to context-expression pairs, then it is also impossible to state in a finite manner the semantic rules for the associated semantics that assigns characters directly to expressions.

---

88 Lasersohn (2014, 186) suggests this kind of rules, but he doubts that natural languages could ever use them. As he rightly points out “allowing rules like this may be threatening to the idea that formal, grammar-based semantics bear most of the *explanatory burden* in the theory of truth conditions, but it is not threatening to the idea of [weakly] compositional interpretation, or to the idea that in principle, contents are assigned as part of grammar” [my emphasis]. This is, in fact, what truth-conditional pragmatists argue for: to share the burden in the derivation of truth-conditions between semantics (that handles what is linguistically encoded in sentences) and pragmatics (that handles other contextual operations not driven by the linguistic meaning of words).

## 7. Domestications.

Defenders of weak compositionality might argue that the problems I raise are not, strictly speaking, shortcomings of the principle of weak compositionality. They arise because the composition function varies freely with the context of utterance. But for sure we can have weakly compositional semantics' in which the modes of combining meanings vary across contexts as a function of some fixed set of contextual parameters. In this section, I'll go over a possible way of domesticating context-sensitive semantic rules. First I'll present the proposed domestication, and after that I'll point out some of its shortcomings.

Here is an example of how such a semantic rule would look like. Consider a theorist who, for whatever reasons, believes that the quantified noun phrase “every student” in (5) varies its content as a function of certain contextual parameter(s) (e.g. the most salient group of individuals at the time and place of utterance) but she believes that neither “every” nor “student” vary their content across contexts of utterance.

(5) Every student failed the exam.

Such a theorist would build the restriction of the domain of quantification into the content of the quantified noun phrase without building it into the content of the constituents: neither in the content of the determiner “every” nor in that of the noun “students” - nor are there, according to this account, any other constituents (i.e. variables) in the syntactic structure of the quantified noun phrase). Then the content of the quantified noun phrase can vary across contexts of utterance without any corresponding variation in the content of its constituents. A theory built along these lines will fail to be strongly compositional, but it *can* be weakly compositional. One way to implement this idea is to have the restriction of the domain of quantification introduced by the semantic operations that determine the content of quantified noun phrases from the content of its constituent determiner and noun phrase. The semantic rule that introduces these operations would, then, have the following form<sup>89</sup>: if *e* is an expression whose immediate constituents are an expression of category Det and an

---

89 Pelletier (2003) and Lasersohn (2012) put forward this kind of semantic rules. Pelletier endorses it as part of the correct analysis of quantified noun phrases, while Lasersohn is more reserved.

expression of category NP, then:

$$I(e,C) = I(e_{\text{Det}},C)( I(e_{\text{NF}},C) \cap RC ) \text{ where } RC \text{ is the set of relevant objects in } C.$$

This rule defines a different composition function for each context. Still, as opposed to the rules considered in the previous section, the rule is finitely storable. A semantics that employs this kind of rules will then employ different composition functions in different contexts and the composition function will vary with the set of contextually relevant objects. Given that the composition function varies with a fixed set of contextual parameters (in this case, with the contextually relevant objects) some of the worries that I raised in the previous section should disappear. For there is a way to formulate and learn the semantic rules that introduce these composition functions that vary with the context of utterance.

But I want to raise two other worries concerning this type of rules. The first thing that should be noted is that this kind of rule undermines much of the motivations for wanting weak compositionality. Rules of this kind will be of no help to the theorist that wants to allow for flexible and rich pragmatic intrusion on truth conditions. In other words, this kind of rules are useful when it comes to contextual variation that is akin to indexicality. Secondly, I'll argue that this kind of rules presuppose an implausible view of the semantic significance of syntactic rules.

How can this rule be used to handle, in a weakly compositional account, all (or most) cases of context-sensitivity put forward by the skeptics? How can this rule be used by truth-conditional pragmatists, for example? Remember truth-conditional pragmatists accept the data put forward by the skeptics, and accept that virtually any declarative sentence is affected by pervasive and creative forms of context sensitivity. They claim that the contextual variation in the truth-conditions of natural language sentences is not plausibly the result of fixing the value of some parameter or other, as it is the case with indexicals. Are rules of this type of any use to the truth-conditional pragmatist? Can this type of rule be used to deal in a weakly compositional format with sentences (6)-(9) and other sentences put forward by the skeptics?

(6) The leaves are green

(7) The ball is round

(8) The cat is on the mat

(9) The ham sandwich stinks

There seem to be little prospects for that. For if truth-conditional pragmatists are right, there are contextual effects their truth conditions that can't be handled by appeal to a fixed set of parameters. That is, if truth-conditional pragmatists are right each of the (6) to (9) sentences vary their truth-conditions along a potentially open-ended number of features, many of them determined by intricate and highly specialized interests and concerns, and not along a fixed set of contextual parameters. Moreover even if truth-conditional pragmatists would accept that (6), for example, varies its truth-conditions along a fixed set of parameters, there are a multitude of other sentences with the same syntactic structure as (6) and each of them could vary their truth conditions along a different parameter. If truth-conditional pragmatists are right and radical context-sensitivity that gives "the possibility of creative uses [of expressions] that extend language to new and previously unimagined applications<sup>90</sup>" is so pervasive that it affects virtually any sentence with the syntactic structure of (6) it is doubtful that any domesticated rule can be formulated. Why? Because it is doubtful that we can identify a set of contextual parameters in terms of which all sentences with the [[NP][VP]] structure vary their truth-conditions. What is relevant for the variation in truth-conditions of (6) seem to be different from what is relevant for the variation in truth-conditions of (8), and both are different from what it seems relevant for the variation in truth-conditions of (9). I strongly suspect that the domesticated context-sensitive semantic rule presented above is tailored to a specific set of data (concerning one type of semantic phenomenon, namely quantifier domain restriction) and that it can't be generalized to cover a larger set of data.

In fact, it is fair to describe a theory that employs these type of domesticated context-sensitive semantic rules as one where pragmatic intrusion into truth-conditions does not occur. A theory of this ilk is one in which on one hand the quantified noun phrase "every student" in (5) varies its content as a function of certain contextual parameter(s) (e.g. the most salient group of students at the time and place of utterance) and on the other

---

90 Bezuidenhout (2002:123)

hand neither “every” nor “student” vary their content across contexts of utterance.

(5) Every student failed the exam.

Such a theory, then, is one where the contextual effects on the interpretation of complexes are no different than indexical resolution: the content of a quantified noun phrases relative to a context is a function of fixed set of contextual parameters. Or to put it in different terms, such rules allow the theorist to specify the character of quantified noun phrases as a non-constant function that determines their content in terms of the relevant objects at the context of utterance. Then there is no need to appeal to pragmatic intrusion. In itself this is not at all problematic, but it does undermine the motivation for weak compositionality. New reasons for wanting weak compositionality for content, over strong compositionality for content are required. Up to now, what motivated the appeal to a weaker form of compositionality for content was the desire to allow for pragmatic intrusion into truth conditions <sup>91</sup>.

## **8. The Semantic Significance of Syntactic Rules**

My second objection concerns the very idea of context-sensitive rules. The concern is that the the idea of context-sensitive rules gets the semantic significance of syntactic rules wrong and it presupposes an implausible account of what is it to know syntactic rules. I won't develop the argument in full detail but only sketch its main lines.

What semantic theories need to explain is speakers' intuition that the truth-values of (5), and other quantified sentences, vary with the context. Then the claim that quantified sentences vary their truth-values because they vary their content across contexts is part of the explanation of these intuitions<sup>92</sup>. Standardly the explanation of intuitions of truth-

---

91 Pagin and Pelletier 2007, Recanati 2010, Lasersohn 2012, explicitly introduce the notion of weak compositionality as a manner of making room for pragmatic intrusion.

92 This is not at all the only available explanation. For example, Stojanovic (2009) puts forward a relativist account of quantifier domain restriction: a domain of quantification is required for determining truth value of quantified sentences, but it is not required for determining their semantic content. Quantified sentences may receive different truth values when evaluated relative to different domains.

values is parceled in two parts. Speakers have the ability to evaluate (5) for truth because on one hand they have the ability to discern under what conditions (5) is true, and on the other hand they have the ability to determine whether those conditions obtain or not in a given occasion of use. The ability to discern truth-conditions belongs to their linguistic competence, and is the proper object of semantic theories, while the ability to determine how the world is at a given occasion of use is, obviously, a non-linguistic ability. How speakers come to know under what conditions (5) is true? The obvious answer is that they know certain properties of (5) on the basis of which they determine its linguistic meaning. Semantic theories, then, explain this linguistic ability by showing how sentences get their linguistic meaning and by assuming that there is a correlation between those aspects of semantic theories establish this and what speakers know when they know the linguistic meaning of sentences.

How does a theory that employs context sensitive rules, like the one introduced above, can explain competence with (5) and other quantified sentences? How does it explain speakers' intuition that the truth-conditions of (5) vary across contexts of utterance? Very briefly the explanation would go along the following lines. The account starts from the fairly standard observation that in the case of quantified sentences the domain of quantification is associated with constituent quantified phrases and not with the sentences themselves. The truth-conditions of (5) vary because the content of the quantified noun phrase "every student" varies across contexts of utterance. Secondly the account claims that the content of "every man" varies from context to context depending on what individuals are included in the domain of quantification. This much is common to all accounts of quantified sentences that trace the variation in the truth-values of (5) to the context-sensitivity of the quantified noun phrase. What sets weakly compositional accounts apart is the claim that the content of "every" and "student" are fixed across contexts and therefore the source of the contextual variation is not in the content of the constituents but in the semantic rule itself. The manner in which the content of constituents combine varies across contexts of utterance. In other words at different contexts, different functions take the content of the constituent determiner and the content of the constituent noun (at that context) and give the content of the noun phrase at that context. Such a semantics assigns context-insensitive linguistic meanings both to the noun and the quantifier (see chapter 2)

but assigns a context-sensitive meaning to (5): a function from contexts to propositions that returns different propositions as a function of the set of contextually relevant objects.

So far, this explains what is it for (5) to have meaning and how its linguistic meaning is determined. To explain linguistic competence with (5) the theory must also explain how speakers come to know its meaning. So what is it that competent speakers know and on the basis of which they can determine what (5) and other quantified sentences mean? The most obvious, and most plausible, answer is that they grasp or know the linguistic meaning of its constituents and its syntactic structure. Because they know the linguistic meaning of the quantified noun phrase “every student” and because its meaning is context-sensitive, they know that (5) varies its truth-conditions across contexts of utterance. We can ask the same question about the quantified noun phrase itself. On the basis of what do competent speakers determine what “every student” means? And on the basis of which they can determine that “every student” denotes different sets of individuals at different contexts? Again, the most plausible and obvious answer would be that they know the meaning of its constituents and they know its syntactic structure. But at this level, knowledge of linguistic meaning of constituents won't explain speakers' knowledge of linguistic meaning of the quantified noun phrase, more exactly of the fact that they know that the quantified noun phrase is context-sensitive. For according to weakly compositional theories, the quantified noun-phrase is context sensitive, while its constituents “every” and “student” are context insensitive (And according to this account, “every” and “student” are the sole constituents of the quantified noun phrase in (5))<sup>93</sup>.

It must be, then, that what explains speakers' intuitions about the variation of truth-conditions of (5) is their knowledge of syntactic structure and its semantic significance. Speakers know that (5) varies its truth-conditions across contexts of utterance, not in virtue of knowing the linguistic meaning of its constituents but in virtue of knowing its syntactic structure. For some expressions (e.g. quantified sentences) it is enough that one knows or grasps their syntactic structure in order to know that its truth-conditions vary across

---

93 I should emphasize that these are claims about linguistic meaning and not about content. An advocate of weak compositionality is still wedded to the idea that linguistic meaning of complex expressions is determined by the linguistic meaning of their constituents and its syntactic structure, and to the idea that knowledge of linguistic meaning of constituents and knowledge of syntax is sufficient to yield knowledge of linguistic meaning of complexes.

contexts of utterance. In other words knowledge of syntactic rules brings about substantive knowledge about actual and possible truth-evaluation of sentences built with that rule. In other words, theories that employ context-sensitive semantic rules attribute the variation in truth-conditions of (5) to the semantic effects of the syntactic rule that combines determiners and nouns, in the sense that at different contexts of utterance different semantic operations combine the meaning of the determiner and the meaning of the noun. Speakers that know the syntactic rule that combines a determiner with a noun into a noun phrase will know, in virtue of that, that sentences that contain quantified noun phrases will vary their truth-conditions across contexts of use.

This does not seem right, though. It does not seem to be the kind of knowledge that one can gain from knowledge of syntax. In fact, such a take on the semantic import of syntactic rules, and on what one knows when she knows syntactic rules goes against most approaches about what knowledge of syntax is. Let me briefly point out why this is so. Intuitions about grammaticality are taken to reliably reflect knowledge of syntactic rules, in the same way in which intuitions about truth-conditions are taken to reliably reflect knowledge of linguistic meaning. In other words, there are good reasons to believe that these different types of intuitions reflect different aspects of linguistic competence<sup>94</sup>.

What is it to say that native speakers know the syntactic structure of expressions, or, equivalently the syntactic rules by which complex expressions are built? Here is a neutral and non-committal way of describing this. Consider some textbook examples (I borrow the examples from Marantz 2005). Native speakers consistently judge (10a), and sentences of similar form, to be grammatically acceptable, while they judge (10b), and sentences of similar form, not to be grammatically acceptable. They also judge (11a), and sentences of similar form, to be grammatically acceptable, but (11b), and sentences of similar form to be unacceptable or incorrect.

(10a) The scared man jumped from his seat.

(10b) \*The man scared jumped from his seat.

---

<sup>94</sup> One important issue concerning knowledge of syntax springs to mind, but I will slur over it as it is not relevant for the present discussion: it is a matter of dispute whether knowledge of syntactic rules is propositional knowledge (or knowledge-that), or some form of practical ability (or knowledge-how). What is relevant is that we can say that competent speakers of a language know the syntactic rules of a language.



(11a) The man scared of porcupines jumped from his seat.

(11b)\*The scared of porcupines man jumped from his seat.

We, then, say that competent speakers know that in English adjectives without complements come before the noun they modify while adjectives with complements come after.

There are some aspects of knowledge of meaning that are explained by speakers' knowledge of syntax. For example, we say that competent speakers know that (3) and (4) differ in meaning because they know that they have different syntactic organizations.

(3) John loves Mary

(4) Mary loves John

We say that competent speakers know that "Mary" occupies different syntactic positions in (3) and in (4), and know the semantic significance of this, namely that the meaning of the name "Mary" combines with the meaning of different verb-phrases in (3) and (4) that occurs in its sister position.

It appears, then, that knowledge of syntax is purely formal, in the sense that it involves only knowledge of relations among expressions and it doesn't involve knowledge of relations between expressions and non-linguistic entities. This is fundamentally different from knowing, for example, the linguistic meaning of sentences (which involves knowing what states of affairs have to obtain in order to make a sentence true). This claim is also supported by the fact that intuitions of grammaticality (which we take to reliably reflect knowledge of syntax) can be divorced from intuitions about truth-conditions (which we take to reliably reflect knowledge of meaning). For example, consider (12a) and (12b):

(12a) Colorless green ideas sleep furiously

(12b) \*Ideas furiously colorless sleep green

Neither of them is meaningful. But native speakers tend to judge the first one as being grammatically correct or grammatically acceptable while the second one as being incorrect or in-acceptable. So native speakers' intuitions about expressions being grammatical or ungrammatical are quite independent of their intuitions about expressions being meaningful

or not. Theories of natural languages have a simple explanation for this: (12a) is syntactically correct in the sense that it is formed according to the syntactic rules of English, while the (12b) is not formed according to syntactic rules of English.

Consider how we explain the meaninglessness of (12a). "Colorless green ideas sleep furiously" is meaningless because we can't determine how the world has to be for it to be true or false. When we compose the referents of its simple constituents what we get is not something against which sentences can possibly be evaluated for truth or falsity. Knowledge of meaning is essentially tied to (actual and possible) truth-evaluations. That is, knowledge of meaning essentially involves correlating linguistic expressions with non-linguistic entities: correlating sentences with states of affairs (or sets of possible worlds) and subsentential expressions with individuals, objects, properties, relations, events and so on. When we combine the referents of constituents of (12a) according to its syntax we don't get anything that can be mapped into truth or falsity. Nevertheless, we know how to combine the meanings of constituents of (12a) even if the result of the combination is, so to speak, vacuous. And we know how to do that because we know how to combine the expressions themselves, that is we know the syntax of (12a) and its semantic significance.

According to advocates of context-sensitive rules, competent speakers know that different states of affairs make (5) true, more exactly at different contexts of use different states of affairs, or different sets of possible worlds, make (5) true. And they know that not not in virtue of knowing the linguistic meaning of its constituents but in virtue of knowing its syntactic structure and its semantic effect. Then, for some sentences, it is enough that one knows their syntactic structure in order to know that at different contexts, different states of affairs make them true. This is quite substantive knowledge about meaning that competent speakers gain simply from knowledge of syntactic rules. As I said, this is an extremely implausible view of what knowledge of syntax (and of its semantic effects) brings in the economy of knowing the linguistic meaning of expressions.

## **9. Summing up**

In this chapter I did two important things. I argued that weak compositionality fails to explain the productive character of natural languages. The reason for this is because weak compositionality allows that the manner in which meanings combine varies freely with the context of utterance. More precisely, the reason why this is in conflict with productive features of languages is because it is impossible to state in a finite manner the semantic rules that introduce the manners in which meanings combine. I, then, discussed a possible domestication by which the ways in which meanings combine do not vary freely with the context of utterance, but vary as a function of certain set of contextual parameters. One important drawback of such domestications is that it undermines the initial motivation for weak compositionality. Finally, I ended by pointing out a serious problem for context-sensitive semantic rules: they get the semantic significance of syntactic rules wrong.



## CHAPTER 6: Conclusions

This essay deals with what constraints different principles of compositionality place on the treatment of context-sensitivity. I focused primarily on whether the claim that natural language sentences exhibit pervasive and radical context sensitivity can be accommodated within a compositional semantics. By pervasive and radical context-sensitivity I have in mind the kind of context-sensitivity that cannot be handled in terms of a fixed set of contextual parameters, and that, allegedly, affects virtually any natural language sentence.

Most authors approach this debate from the perspective set by the following questions. Is the claim that natural language sentences exhibit this kind of context-sensitivity well supported? Are the arguments that supposedly establish the existence of such context-sensitivity good enough? More precisely, are the intuitions raised by skeptics, and their companions the truth-conditional pragmatists, data that semantic theories need to explain? And secondly, is it the case that no semantic theory done in the framework introduced in chapter one can properly handle this kind of context-sensitivity? Authors that are sympathetic towards standard semantic theories try to show that there is no good case for pervasive and radical context-sensitivity. They either try to show that the intuitions put forward by skeptics are not data that semantic theories must account for, or they try to show that standard semantic theories can accommodate these intuitions – i.e. they try to show that this context sensitivity is not that radical.

I approached this debate from a slightly different angle. I didn't weigh the arguments in favor or against the claim for pervasive and radical context-sensitivity. I didn't, for example weigh in on whether the intuitions put forward by skeptics must be predicted by semantic theories. I merely conceded that most of them must be. Nor did I weigh in on whether any of the proposals to account for these intuitions is successful. (In chapter three I

briefly sketched the general lines of the responses available to a theorist working within this framework). The angle from which I approached the dispute is the following: assuming (or conceding) that there is pervasive radical context-sensitivity, is this incompatible with the claim that natural language semantics are compositional? To give a precise answer to this question I distinguished three principles of compositionality, where each of them place different constraints on semantic theories. I showed in chapter three that if we accept that there is pervasive and radical context sensitivity this is incompatible with the notion of strong compositionality but not with that of weak compositionality nor with that of character compositionality.

The question I addressed can be formulated in the following way: Is the contribution of the context to the determination of truth-conditions of sentences restricted to that of fixing the values of context-sensitive expressions (or other material) present in the logical form? A positive answer that, yes the only role that context can play is to fix the value of context-sensitive material realized in the logical form of sentences, is compatible with a strongly compositional semantics. A negative answer, though, is incompatible with a strongly compositional semantics, but is compatible with a weakly compositional or a character compositional semantics.

In fact, the very idea of weakening the principle of compositionality for content was motivated by the the desire to have a semantics in which the context plays a richer role in the determination of truth conditions and that satisfies “some form of compositionality”. Of course, we might ask why are these theorists interested that *some form* of compositionality is satisfied? The answer is that compositionality is desirable because it explains certain properties of natural languages and of linguistic competence with them. In chapters four and five I argued that weak compositionality fails to deliver some of these expected benefits: it fails to deliver systematicity and it fails to explain the productive features of natural languages. Weak compositionality allows context to contribute in an indefinite number of ways to the truth-conditions of sentences over and above its contribution to the interpretation of their simple parts. The conflict comes from the fact that it is impossible to formulate in a finite manner the semantic rules for a semantics that allows this role for the context. The semantic rules will introduce for every context a distinct semantic operation,

which means that they will have to take the form of indefinitely long disjunctions. Such rules are obviously not finitely statable nor learnable. If context-sensitive semantic rules are properly domesticated, in the way discussed in chapter five, this will ease off the pressure of not explaining the productive features of natural languages. But, even together with these domestications, context-sensitive semantic rules presuppose an unrealistic picture of the semantic significance of syntactic rules: knowledge of syntactic rules alone will provide knowledge of variance in truth-evaluations. I believe that all these are good reasons to doubt the viability of (compositional) semantic theories that allow the contribution of the context to go beyond fixing the content of simple expressions (or variables) present in the logical form of sentences.

If the contribution of the context to the determination of truth-conditions of sentences at contexts restricted to that of fixing the values of context-sensitive expressions (or other syntactic material) present in the logical form then what Szabo calls the context-thesis is also true: the content of a complex expression depends on the context only insofar as the content of its simple constituents do. This offers a good reason to prefer strongly compositional semantics' over weakly compositional ones. Results of chapter two are relevant here. First, the context-thesis is entailed by strong compositionality, but does not entail strong compositionality. So that a semantics satisfies the context-thesis doesn't mean that it satisfies strong compositionality. It might fail to be strongly compositional for other reasons - e.g. if substitution of synonyms is not meaning-preserving in that semantics. Secondly, a semantics can satisfy the context-thesis and satisfy weak-compositionality. But this will be a semantics in which the extra contribution of the context will always be vacuous or it will be constant. Moreover such a semantics will, in fact, also satisfy strong compositionality. But we shouldn't take this claim about preferences too restrictively. How about character compositionality (compositionality of linguistic meaning)? After all, according to results of chapter two, any weakly compositional semantics is also character compositional.

Character compositionality cuts both ways. If the context-thesis (about content) is true, then obviously also the following thesis is true, (let's call it the context-thesis for linguistic meaning): a complex expression is context-sensitive only insofar as some of its

simple constituents are context-sensitive. If a complex expression has a context-sensitive linguistic meaning, this context-sensitivity is traceable to some of its simple constituents. A semantics that fails to satisfy the context-thesis, can, nevertheless, satisfy character compositionality. But the considerations from chapter five apply here as well. For example, a semantics for quantified sentences that employs the type of domesticated semantic rule discussed in chapter five is a semantics that satisfies character compositionality but fails to satisfy the context-thesis. In such a semantics the linguistic meaning of quantified sentences will be represented as non-constant functions from contexts to intentions and the linguistic meaning of their simple constituents can be represented as constant functions from context to intentions. In principle, even properly character compositional semantics can fail to satisfy the context-thesis - by a properly character-compositional semantics I mean a semantics that satisfies character compositionality but fails to satisfy weak compositionality. (For example, semantics that contain context-shifting operators are properly character compositional.)

Once we leave the level of strong compositionality, considerations about compositionality alone do not impose adherence to the context thesis. One has to look for other motivations. I believe that considerations about the form that semantic rules must take if the context-thesis does not hold, can offer good motivations for adherence to the context-thesis.

If the arguments from chapters four and five are correct then we want accept some form of the context-thesis (either about content or about linguistic meaning). The options to deal with the intuitions put forward by skeptics are, then, identical with the ones described in chapter three. The first, is to claim that these intuitions are not about the truth-conditions of sentences but about the pragmatics of their utterances. The second is to claim that intuitions about the variations in truth values can be explained by the semantics of some of the constituents of these sentences: some constituents are either context-sensitive or they are circumstance sensitive.



## References

- Bach, Emmon (1988) *Categorial Grammars as Theories of Language*. In Oehrle, R. Bach, E. and Wheeler, D. (eds.) (1988) *Categorial grammars and natural language structures*. Dordrecht: Reidel: 17-34.
- Bach, Emmon (1989) *Informal Lectures on Formal Semantics*, Albany, State University of New York Press.
- Bach, Kent. (2002) *Seemingly Semantic Intuitions*. In J. Campbell, M. O'Rourke, and D. Shier, eds., *Meaning and Truth: Investigations in Philosophical Semantics*. New York: Seven Bridges Press: 284-292.
- Bach, Kent. (2012). *Context-Dependence (Such as it is)*. In Garcia-Carpintero, M. and Kolbel, M. (eds). *The Continuum Companion to the Philosophy of Language.*, London – New York, Continuum Publishing.
- Bach, Kent & Bezuidenhout, Anne (2002) *Distinguishing Semantics and Pragmatics*, in Joseph K. Campbell, Michael O'Rourke & David Shier (eds.) *Meaning and Truth: Investigations in Philosophical Semantics*. Seven Bridges Press. 284-310.
- Berg, J. (2002) *Is Semantics Still Possible?*, *Journal of Pragmatics* 34.3, 349–359.
- Bezuidenhout, A.: (1997), *Pragmatics, semantic underdetermination and the referential/attributive distinction*, *Mind* 106, 375–409.
- Bezuidenhout, Anne. (2002). *Truth-conditional Pragmatics*, *Philosophical Perspectives*, 16: 105-34.
- Borg, Emma (2004a) *Formal Semantics and Intentional States*, *Analysis* 64.3, July 2004, pp. 215–23.
- Borg, E. (2004b). *Minimal Semantics*. Clarendon Press, Oxford.
- Cappelen, Herman, & Lepore, Ernie (2005). *Insensitive semantics*. Oxford: Blackwell.
- Carnap Rudolf (1928/2003), *Pseudoproblems in Philosophy*, reprinted in *The Logical Structure of the World and Pseudoproblems in Philosophy*. Open Court.
- Carnap, Rudolf (1947), *Meaning and necessity*, (first edition) Chicago: University of Chicago

- Press.
- Carston, Robbin. (2002), *Thoughts and utterances*, Oxford: Blackwell.
- Chomsky, Noam (1957), *Syntactic Structures*, Mouton.
- Chomsky, N (1965), *Aspects of the Theory of Syntax*, Cambridge, Mass. MIT Press,
- Chomsky, N (2000), *New Horizons in the Study of Language and Mind*, Cambridge University Press.
- Cohen, L. Jonathan (1985), A Problem about Ambiguity in Truth-Theoretical Semantics, *Analysis* 45 (3):129 - 134.
- Cohen, L. Jonathan (1986), How is conceptual innovation possible?, *Erkenntnis* 25 (2):221 – 238.
- Elbourne, Paul (2008), The argument from binding, *Philosophical Perspectives* 22 (1):89-110.
- Davidson, D.(1965), Theories of meaning and learnable languages, Reprinted in D. Davidson, *Inquiries into Truth and Interpretation*, Oxford: Clarendon Press, 2001: 3–16.
- Davis, Wayne A. (2003), *Meaning, Expression, and Thought*, Cambridge University Press.
- Dever, J. (1999). Compositionality as methodology, In *Linguistics and Philosophy*, 22: 311–326.
- Dever, J.(2006). Compositionality. In Ernest Lepore & Barry Smith (eds.), *The Oxford Handbook of Philosophy of Language*. Oxford University Press. 633—666.
- Dowty, David (2007). Compositionality as an empirical problem. In Chris Barker & Pauline I. Jacobson (eds.), *Direct Compositionality*. Oxford University Press. 14--23.
- Fodor, J. (1983). *Modularity of Mind*. Cambridge, Mass.: MIT Press.
- Fodor, Jerry, & Lepore, Ernest (1991) *The compositionality papers*. Oxford: Blackwell, 2002.
- Fodor, J. (2001). Language, thought and compositionality. *Mind and Language*, 16, 1–15.
- Fodor, J. & Lepore, E. (2004). Out of Context. *Proceedings and Addresses of the American Philosophical Association* 78 (2):77 – 94.
- Grandy, R. (1990). Understanding and the principle of compositionality. *Philosophical Perspectives* 4:557-572.
- Grice, P. (1989) *Studies in the ways of words*. Cambridge Mass. Harvard University Press
- Groenendijk, J. & Stokhof, M. (2005). Why Compositionality?, In Greg N. Carlson & Francis Jeffrey Pelletier (eds.), *Reference and Quantification: The Partee Effect*. CSLI Publications. 83—106.

- Heim, Irene & Kratzer, Angelika (1998). *Semantics in Generative Grammar*. Blackwell.
- Hodges, W. (2001) Formal features of compositionality. *Journal of Logic, Language, and Information*, 10: 7–28.
- Hodges, W. (2014). Compositionality is not the problem. *Logic and Logical Philosophy* 6:7.
- Jackendoff, R. (1992) Mme. Tussaud Meets the Binding Theory in *Natural Language & Linguistic Theory*, Vol. 10, No. 1 (Feb., 1992), pp. 1-31.
- Jackendoff, R. (1997). *The Architecture of the Language Faculty*. Cambridge, Mass. MIT Press.
- Janssen, T. M. V. (1983). *Foundations and Applications of Montague Grammar*. Amsterdam: Mathematisch Centrum.
- Janssen, T. M. V. (1997) Compositionality. in J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, Amsterdam: Elsevier. 417–473.
- Janssen, T. M. V. (2001). Frege, contextuality and compositionality. *Journal of Logic, Language and Information*, 10, 115–136.
- Kaplan, D. (1977). Demonstratives. An essay on the logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Perry & H. K. Wettstein (Eds.), *Themes from Kaplan*. Oxford: Oxford University Press, 1989.
- King, Jeffrey (1996). Structured propositions and sentence structure. *Journal of Philosophical Logic* 25 (5):495 – 521.
- King, J., & Stanley, J. (2005). Semantics, pragmatics, and the role of semantic content. Reprinted in Stanley, J. (2007) *Language in Context: Selected Essays*. Oxford, Oxford University Press.
- Kölbel, M. (2004) Faultless disagreement. *Proceedings of the Aristotelian Society*, 104(1), 53–73.
- Kölbel, M (2008) The evidence for relativism. *Synthese* 166 (2):375-395.
- Kölbel, M (2008) Motivations for relativism. In Manuel García-Carpintero & Max Kölbel (eds.), *Relative Truth*, Oxford, Oxford University Press.
- Magidor, Ofra (2009). Category mistakes are meaningful. *Linguistics and Philosophy* 32 (6):553-581.
- Nunberg, Geoffrey (1995). Transfers of Meaning. *Journal of Semantics* 12 (2):109-132.

- Lahav, R. (1989). Against compositionality: The case of adjectives. *Philosophical Studies*, 57, 261–279. Reprinted in Byrne, Darragh & Kölbel, Max (eds.) (2010). *Arguing About Language*. Routledge.
- Laserson, P. (2012) Contextualism and Compositionality. *Linguistics and Philosophy*, 35 (2):171–189.
- Lewis, D. (1970). General semantics. Reprinted in Davis, S., and Gillon, B., Brendan, S. (eds.) *Semantics: A Reader*. Oxford University Press.
- Lewis, David (1975) Language and Languages. In Keith Gunderson (ed.) *Minnesota Studies in The Philosophy of Science*. University of Minnesota Press. 3-35.
- Lewis, D. (1979), Scorekeeping in a Language Game, *Journal of Philosophical Logic* 8, 339–359. Reprinted in D. Lewis, *Philosophical Papers, Volume 1*, Oxford University Press, Oxford, 1983.
- Lewis, D. (1980). Index, Context, and Content. in Kanger, S. and Ohman, S. (eds.), *Philosophy and Grammar*, Dordrecht: Reidel: 79–100; reprinted in Lewis, D. 1998. *Papers in Philosophical Logic*. Cambridge: Cambridge University Press.
- Pagin, P. (1997) Is compositionality compatible with holism?. *Mind & Language*. 12:11-33.
- Pagin, P. (2003). Communication and strong compositionality. *Journal of Philosophical Logic* 32, pp. 287 - 322.
- Pagin, P. (2005). Compositionality and context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy: Knowledge, meaning and truth* (pp. 303–348). Oxford: Oxford University Press.
- Pagin, P., & Pelletier, J. (2007). Context, content and communication. In G. Preyer & G. Peter (Eds.), *Context-sensitivity and semantic minimalism* (pp. 25–62). Oxford: Clarendon Press.
- Pagin, P & Westerståhl, D. (2010a) Compositionality I: Definitions and Variants, in *Philosophy Compass* 5:265-82, 2010
- Pagin, P & Westerståhl, D. (2010b) Compositionality II: Arguments and Problems, *Philosophy Compass* 5:250-64, 2010.
- Partee, Barbara (1979) Constraining Transformational Montague Grammar: A Framework and a Fragment. In Davis, S. and Mithun, M (eds), *Linguistics, Philosophy and Montague Grammar*, Austin, University of Texas Press.

- Partee, Barbara H. (1995). Quantificational structures and compositionality. In Emmon Bach, Eloise Jelinek, Angelika Kratzer & Barbara Partee (eds.), *Quantification in Natural Languages*. Kluwer. 541—6.
- Partee, Barbara H. (1997). Montague semantics. In J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, Amsterdam: Elsevier 7-82.
- Partee, Barbara H. (1997). Genitives: A case study. In J. van Benthem and A. ter Meulen (eds.), *Handbook of Logic and Language*, Amsterdam: Elsevier 464-470.
- Partee, Barbara (2004). *Compositionality in Formal Semantics: Selected Papers of Barbara Partee*. Blackwell Pub.
- Pelletier, F. J. (1994). The principle of semantic compositionality. *Topoi*, 13, 11–24. expanded version reprinted in Davis, S., and Gillon, B., Brendan, S. (eds.) *Semantics: A Reader*. Oxford University Press: 133-157.
- Pelletier, F. J. (2003). Context-dependence and compositionality. *Mind and Language*, 18, 148–161.
- Perry, J. (1986). Thought without representation. In *Proceedings of the Aristotelian Society* (Suppl. Vol. 60, pp. 263–283). London: Aristotelian Society Publications.
- Perry, J. and Israel, D (1996) Where Monsters Dwell. In Jerry Seligman & Dag Westerståhl (eds.) *Logic, Language and Computation*, CSLI Publications, Stanford: 303 – 316.
- Pullum, Geoffrey K. and Barbara C. Scholz (2010), Recursion and the infinitude claim. In Harry van der Hulst (ed.), *Recursion in Human Language* (Studies in Generative Grammar 104),. Berlin: Mouton de Gruyter: 113-138.
- Predelli, S. (2005a), *Contexts. Meaning, Truth, and the Use of Language*. Oxford University Press.
- Predelli, S. (2005b). Painted Leaves, Context, and Semantic Analysis. *Linguistics and Philosophy* 28.3: 351-74.
- Recanati, F. (1995). The alleged priority of literal interpretation. *Cognitive Science*, 19, 207–232.
- Recanati, F. (2001). What is said. *Synthese*, 128, 75–91.
- Recanati, F. (2002). Unarticulated constituents. *Linguistics and Philosophy* 25: 299–345.
- Recanati, F. (2002). Does linguistic communication rest on inference? *Mind & Language*, 17, 105–126.

- Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- Recanati, F. (2010). *Truth-conditional pragmatics*. Oxford: Oxford University Press.
- Reimer, Marga (2002). Do Adjectives Conform to Compositionality?, *Noûs* 36 (s16):183 - 198.
- Rothschild, D. & Segal, G. (2009), Indexical Predicates. *Mind and Language* 24 (4):467-493.
- Sainsbury, R. M. (2001). Two ways to smoke a cigarette. *Ratio* 14 (4):386–406.
- Schaffer, Jonathan (2012). Necessitarian propositions. *Synthese* 189 (1):119-162.
- Schlenker, Philippe (2003). A plea for monsters. *Linguistics and Philosophy* 26 (1):29-120.
- Searle, J. R. (1978). Literal meaning. *Erkenntnis*, 13, 207–224.
- Searle, J. (1980). The background of meaning. In J. Searle, F. Kiefer, & M. Bierwisch (Eds.), *Speech act theory and pragmatics* (pp. 221–232). Dordrecht: Reidel.
- Searle, John (1992), *The Rediscovery of the Mind*, Cambridge, Mass. MIT Press.
- Soames, Scott (1984). Linguistics and psychology. *Linguistics and Philosophy* 7 (2):155 – 179.
- Soames, Scott (1985). Semantics and psychology. In Jerrold J. Katz (ed.), *The Philosophy of Linguistics*. Oxford University Press. 204--226.
- Soames, Scott (1989). Semantics and semantic competence. *Philosophical Perspectives* 3:575-596.
- Soames, Scott. (2005). Naming and asserting. In Z. G. Szabó (Ed.), *Semantics vs. Pragmatics* Oxford: Oxford University Press: 356–382.
- Soames, Scott (2012). *What is Meaning?*. Princeton University Press.
- Stalnaker, R. (1998). On the representation of context. *Journal of Logic, Language and Information* 7 (1):3-19. reprinted in Stalnaker, R (1999). *Context and Content: Essays on Intentionality in Speech and Thought*, Oxford, Oxford University Press.
- Stanley, J. (2007) *Language in Context: Selected Essays*. Oxford, Oxford University Press.
- Stanley, J., & Szabó, Z. G. (2000). On quantifier domain restriction. *Mind and Language*, 15, 219–261.
- Szabó, Z. G. (1995). *Problems of compositionality*. PhD Thesis, MIT.
- Szabó, Z. G. (2000a). Compositionality as supervenience. *Linguistics and Philosophy*, 23, 475–505.
- Szabó, Z. G. (2000b/2010). Adjectives in context. In I. Kenesei & R. Harnish (Eds.), *Perspectives on semantics, pragmatics, and discourse* (pp. 119–146). Amsterdam: John Benjamins, reprinted in Byrne, Darragh & Kölbel, Max (eds.) (2010). *Arguing About*

- Language*. Routledge.
- Szabó, Z. G. (2008). Structure and conventions. *Philosophical Studies* 137 (3):399 – 408.
- Szabo, Z. (2010). The determination of content. *Philosophical Studies* 148 (2):253 - 272.
- Szabó, Z.G. (2013) Compositionality, *The Stanford Encyclopedia of Philosophy* , E. N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/fall2013/entries/compositionality/>>.
- Travis, C. (1985). On what is strictly speaking true. *Canadian Journal of Philosophy*, 15, 187–229.
- Travis, C. (1994). On constraints of generality. In *Proceedings of the Aristotelian Society* (New Series 44, pp. 165–188). London: Aristotelian Society Publications.
- Travis, C. (1996) Meaning's Role in Truth, *Mind* 105, 451–466.
- Travis, C. (1997). Pragmatics. In B. Hale & C. Wright (Eds.), *A companion to the philosophy of language* . Oxford Blackwell: 87–107.
- Travis, C. (2008). *Occasion-sensitivity: Selected essays*. Oxford: Oxford University Press.
- Weiskopf, D. A. (2007). Compound nominals, context, and compositionality. *Synthese*, 156, 161–204.
- Westerståhl, D (1998) On Mathematical Proofs of the Vacuity of Compositionality, In *Linguistics and Philosophy*, vol. 21, 6, 635-643.
- Westerståhl, D. (2004). On the compositional extension problem. *Journal of Philosophical Logic* 33 (6):549-582.
- Westerståhl, D. (2012). Compositionality in Kaplan style semantics. In M. Werning, W. Hinzen, & E. Machery (Eds.), *The Oxford Handbook of compositionality* (pp. 192–219). Oxford: Oxford University Press.
- Ziff, P. (1972). *Understanding understanding*, Cornell University Press.