Semantic Compositionality

Francis Jeffry Pelletier

Contents

Semantic Compositionality[1]

Francis Jeffry Pelletier
Department of Philosophy
University of Alberta, Canada
jeff.pelletier@ualberta.ca

Keywords: compositionality, meaning, holism, semantics, Frege, context

1 Summary

Most linguists have heard of semantic compositionality. Some will have heard that it is the fundamental truth of semantics. Others will have been told that it is so thoroughly and completely wrong that it is astonishing that it is still being taught. The present chapter attempts to explain all this. Much of the discussion of semantic compositionality takes place in three arenas that are rather insulated from one another: (a) philosophy of mind and language, (b) formal semantics, and (c) cognitive linguistics and cognitive psychology. A truly comprehensive overview of the writings in all these areas is not possible here. However, this entry does discuss some of the work that occurs in each ofl these areas. A bibliography of general works, and some internet resources, will help guide the reader to some further, undiscussed works (including further material in all three categories).

2 Functional Compositionality

Semantic Compositionality is (usually) defined thus:

Definition 1. (Semantic Compositionality) *The meaning of a syntactically complex unit is a function of the meanings of the complex unit's simpler parts and the syntactic way those parts are put together.*

This is also called *functional* semantic compositionality, or simply functional compositionality, due to the appearance of the word 'function' in the definition. (I will also call it 'functional compositionality', when I wish to separate it from other conceptions of compositionality.)

There are various other notions of compositionality in the literature, some nearly identical with this one while others are rather different in one way or another. We will encounter some of these in §§7 and 8 below.

As well, one can search for an underlying metaphysical or epistemological theme that is somehow behind Definition (1), and once such a theme is identified, one can see how some slight modifications of (1) could result in theories of meaning that are still within the theme, but not quite in accord with the original definition of semantic compositionality. In turn, this can result in rival accounts of the data that are used to support or justify Definition (1), because these rival accounts nonetheless share the same background theme. That is to say, semantic compositionality of the functional sort is only one way to carry out the program that the background theme presumes. I will discuss this background theme in §7.1 (where it is identified as Atomism), and follow it with examples in §8.4 and §9.1.2 below.

And of course, corresponding to all these variations on Definition (1), one can find theories that are opposed to them. These opposed theories often go by the name *semantic holism* – or in those cases where we look to the background theories, they are sometimes called *wholism*[2] – although the holism/wholism names cover a rather wide variety of theories that may have no real similarity to one another, except in denying (1), or denying some theme that is background to Definition (1).

This sort of excavation of background assumptions and taxonomy of theories is not a part of direct semantic theorizing (it is more of a philosophical meta-theoretic analysis of types of linguistic theories), but I think that understanding all this meta-semantic rumination will help in following current semantic theories and their adoption or rejection of semantic compositionality. Then too, semantic compositionality is itself not a semantic theory but rather some sort of property of some semantic theories. Therefore, it seems correct to characterize it as a *meta-semantic* property and to claim that the different conceptions of semantic compositionality, as well as the theories that oppose compositionality, correspond to different semantic meta-theories.

We start by mentioning a very brief early history of semantic compositionality in §3, and follow this in §4 with remarks on some of the concepts that are presumed by Definition (1). At that point we will be in a position to appreciate a few aspects of Definition (1) that might give pause to anyone wondering whether semantic compositionality is a plausible position to hold (§5). We also look (§6) to a number of empirical features of natural language that have been thought to show that semantic compositionality is not true of natural languages. These include issues involving context, idiomatic, and quotation – topics which are usually thought to be the most significant challenges to semantic compositionality, even by those who otherwise would happily adopt semantic compositionality – and I mention approaches taken by some strong advocates of semantic compositionality that perhaps manage to evade the apparent problems. After that, we broaden our viewpoint and consider both

the background presumptions of semantic compositionality (§7) and some alternatives to functional compositionality that are in the literature (§§8–9). This involves considering both the typical arguments in favour of semantic compositionality (§10), and also problems with these arguments. In §11 we look to some of the interesting formal properties are relevant to compositionality and others that show what semantic compositionality can do (or is alleged to be able to do). We close by citing some other general overview works on Semantic Compositionality and also some material available on the web about compositionality (§§13–14).

3 Some Historical Remarks

The idea behind the definition of Semantic Compositionality in Definition (1) is often said to be due to Gottlob Frege (1848–1925), although this attribution of semantic compositionality to Frege is a matter of scholarly dispute (see, for example, Pelletier (2001), Janssen (2001)). As well, various scholars have claimed to find versions of the principle or statements of a similar sentiment in earlier works: for example, Gillon (2007) cites Pāṇini as a very early example; Hodges (2006) has pointed to medieval Arabic scholars and as well to (Husserl 1900: Lecture 4) as a possible source of Frege's thought on this topic. Dever (2006: p. 634ff) cites Lambert of Auxerre; Pagin & Westerståhl (2010a) cite Abelard, Buridan, and Hobbes as having views similar to semantic compositionality. A nice description of the concepts of semantic compositionality and semantic contextuality as they appeared Frege and in the literature contemporaneous with Frege is in Janssen (2012).

The actual term 'compositionality' seems to have first been used in the modern sense by Katz & Fodor (1963), although their interests appear to be more about some sort of "lexical composition". The surge of current interest in the notion can be traced to Partee (1984), which remains a major source of insight into the concept as well as a source for noting difficulties with its application to natural language. After the appearance of Partee (1984) there were numerous publications either defending or attacking semantic compositionality, in each of the linguistic, philosophical, and psychological literatures; and there was a very important addition to the literature by Hodges (1998, 2001) which was elaborated in some later works (Hodges 2012a,b). These works by Hodges have placed a new layer of mathematical sophistication to discussions of semantic compositionality. And of course, both Partee and Hodges acknowledge Richard Montague's inspiration, especially as from Montague (1970: §3).

4 Some Background Concepts Used by Functional Compositionality

Many writers have remarked – usually unfavourably – on the fact that semantic compositionality as defined in (1) is often given without any further explanation of the notions used in the definition: especially *function, meaning,* and *put together*. Different thoughts about these terms can (and have) given rise to different denotations of the term 'semantic compositionality'. As well, there are other definitions and conceptions of compositionality that have been aired and have sometimes occasioned disputes about the truth or plausibility of compositionality, where the dispute really depends on which conception "really" is semantic compositionality.

4.1 "part"

Even though the notion of *part* is determined by the background syntactic theory, there is still a distinction to be made between (what is usually called) *an immediate part* and a part *tout court*. Once again, it is the syntactic theory that determines what is an immediate part of some syntactically-identified unit. But one can also talk about the parts of these immediate parts, and *their* parts, and so on, down to the basic subparts that the syntax recognizes. We shall see in §8.2 that this difference really does make a difference, in characterizing distinct forms of compositionality.

There is also a distinction to be made between a conception of the syntax as embodying parts considered as being characterized by syntactic categories (such as NP, VP, CNP, and so on), a conception where the parts are "logical-form parts", and a conception where we consider only the *string*. Some further remarks about these differences are in §8.3.

4.2 "function"

The meaning of 'function' in the definition of functional compositionality is usually taken to be the mathematical usage, so that it can be conceptualized as a rule which, when presented with certain input information, not only yields some value but, being a *function*, it places a restriction on the rule so that *whenever* the rule is presented with *the same input information*, it will yield the same value. Consider the simple sentence-describing syntactic rule $S \rightarrow NP\ VP$. Functional compositionality here would entail that each specific instance of any particular sentence of that form must have the same meaning, whenever the meanings of the NPs and a VPs are the same.

This is an important point to keep in mind, as some objections to semantic compositionality turn on a different apprehension of what a function is.

4.3 "meaning"

Although the notion of "meaning" has been left unspecified in the definition of functional compositionality, that may not matter for many theories of meaning. (In §9 we will consider notions of "semantic holism" and "subjective content", where the account of meaning makes it difficult to see how the language could be semantically compositional.) Functional compositionality is supposed to be a general constraint on the relation of syntax to semantics and not on the particular syntactic theory or semantic theory under discussion – that is to say, it is a *meta-semantic* comment about how the syntax-semantic interface should be conceived. The idea is that so long as one has some fairly well-specified syntactic theory and some at least partially specified semantic theory, it can be evaluated whether functional compositionality characterizes that pair – at least for the portions that are given. This could make it be that, under some conceptions of meaning for language $\boldsymbol{L}$, semantic compositionality holds, while for other conceptions it doesn't hold. On the reverse side of the question, there can be an issue of whether, when presented with a particular conception of meaning, there is *some* syntactic treatment or other of a given language (seen as a set of strings) that would allow for a semantically compositional semantics. A further and separate issue is whether, from a partially-specified pair of language and compositional semantic theory, it can be extended to a functionally compositional characterization of the full language. We consider this topic below in §11.3.

4.4 "put together"

Of course, for semantic compositionality, the background presupposition is that the "objects" under consideration are syntactic entities or constructions. And that these are "put together" in accordance with a group of syntactic "rules" (or, if you prefer, "descriptions" or "constructions" or . . .). But also, as a part of this idea, each of these ways of "putting together" is associated with some meaning-transforming function that describes how the meanings of the parts – which are the inputs to this function – are altered to form the meaning of the whole. If there is a syntactic rule/description such as $CNP \rightarrow Adj\ N$, then associated with this is a function that describes how to employ the meaning of the adjective and the meaning of the noun in order to uniquely form the meaning of the common noun phrase.

Once the idea is put like this, it has seemed apposite to many authors of semantic textbooks to formulate a grammar in such a way that each syntactic rule is paired with the correlated semantic function. And so the grammar of languages consists of not just the set of syntactic rules, but rather as a set of *pairs* of syntactic-semantic rules. (Often called the "Rule-to-Rule" format of Bach (1976).)

5 Some Theory-Internal Problems for Functional Compositionality

5.1 Functions

To many, the use of the mathematical notion of a function, as described briefly in §4.2 above, doesn't seem right. In fact there are some aspects of the definition that might give pause to anyone wondering whether it is a plausible position to hold. Can it really be right that *every* instance of any one of

(1) a. The sun is shining
    b. The President is happy
    c. Jerry is tall

means the same as any other instance of that sentence? For instance, isn't (1-a) saying one thing when spoken in Minneapolis at local noon on Christmas Day 2001, but saying something different when spoken at 0500 hours in St. Andrews, Scotland, on Eid al-Fitr in 2015? And isn't (1-b) true or false depending on when it is uttered (and truth vs. falsity of instances is sufficient to make a meaning difference, isn't it?). Isn't it also dependent upon *where* such a sentence is uttered: the USA or Argentina, for example, where the presidents are different? Don't utterances of (1-c) differ in truth-value (and hence in meaning?) when they are asserted in the Pimlico jockey dressing room or the Golden State Warriors dressing room?

The options open to supporters of functional compositionality seem to be limited to saying that either the NP subjects of different instances of the same sentence have different meanings, or their VP predicates have different meanings, or that the different instances were syntactically composed by different rules. None of these options seem particularly attractive, and we will later address the sorts of responses that supporters of the functional interpretation of compositionality have given (§6.1).

5.2 Depends On

The apparent task of those who advocate functional compositionality as a goal of semantic theorizing is to find some function or other which will use an already-given syntax and thereby demonstrate the correctness of semantic compositionality for that grammar. Or alternatively, to find where the syntax somehow blocks the "natural meaning function" and thus suggests a place where the syntactic rules need alteration.

But finding some arbitrary function to do this job is easy – too easy, according to many theorists. For, it seems possible, without any further constraints, to gerrymander the

function so as to arbitrarily (but in accord with the function's definition) admit or omit anything at all as the value of any object in the function's domain. If so, then the notion of semantic compositionality would be trivial, or vacuous. In response to this worry about the laxity in the notion of a function, Szabó (2012: p.71) argues for the inclusion of a requirement that the function be one that makes the meanings of the parts *determine* the meaning of the syntactic whole. Pelletier (2012: pp.160–1) says that the function has to correspond to the way that the meaning of the complexes *really depend* on the meanings of their parts... where 'really depend' is taken in some sort of ontologically-explanatory way. Pagin (2009) suggests that the function has to be *a computable function*, or perhaps an "elementarily computable" function – or even a "humanly-computable-by-any-normal-person" function.

Not only could one gerrymander the function, but there are also some "peculiarities" implicit or hidden in the notion of functional compositionality. For instance, if the language happened to have no synonymous items at all – no synonymous lexical items, no synonymous complex items – then by definition it would be semantically compositional. For, to be a counterexample to functional compositionality there would have to be two synonymous items such that the appearance of one of them into a more complex phrase generated a meaning different from what happens when the other expression is in that otherwise identical phrase (which was generated by identical syntactic rules/descriptions).. Without any synonyms there could never be a counterexample. Now, although this possibility doesn't show that compositionality must be wrong, it does seem to go against the general tenor of semantic compositionality.

Furthermore, the unmodified conception of a function allows for there to be completely unrelated meanings assigned to similar, although not completely identical, structures – since the only relevant restriction forced by (mathematical) functionality is identity. Consider

(2) a. My house is large
    b. My house is yellow

Even assuming that these have the same syntactic structure, we could nevertheless concoct the meaning function in such a way that their meanings are (using small caps to represent meanings in an intuitive manner), respectively

(3) a. MY HOUSE IS LARGE
    b. MY TURTLE IS YELLOW

(2-a) and (2-b) have the same subject NP, but are nevertheless this similarity nowhere shows up in their meanings. However, there is a function that will turn this trick: it needs only to specify that the meaning of *large* (or maybe more generally, the meaning of size adjectives) operates "normally" whereas the meaning of *yellow* (or maybe the meaning of

colour terms in general) always picks out the favourite pet of the speaker as the subject of predication. Even though the lack of uniformity can easily be detected in these simple examples, one can imagine more intricate combinations of the similar strategy which would make this counterintuitive effect be less obvious and perhaps prevent straightforward detection. Yet would would want to say that this general type of function goes against the intent of semantic compositionality. The challenge is to say exactly and precisely *how* it goes against it. For, nothing in the *definition* of a function will prevent this. The parallel effect for predicates could be likewise "legal" unless there were some explicit prohibition that is stated for the meaning function:

(4) a. My house is large
    b. My car is large

could be assigned these meanings, respectively:

(5) a. MY HOUSE IS LARGE
    b. MY CAR IS POWERFUL

One might note that although the individual lexical(-ish) items *my house* and *large*, as used in (2) and (4) respectively, can be unambiguous, their meanings *in the relevant sentences that are generated* have different meanings because of how the VP-meaning interacts with the NP-meaning. There is nothing here that is literally in conflict with the principle of semantic compositionality and yet most people would find this interpretation of "depends on the meanings of its syntactic parts and the way the parts are put together" to be out of keeping with the intent of semantic compositionality.

Some (e.g., Johnson (2014)) have thought that this laxness in the definition of functional compositionality tells against the "pure mathematical notion" of functionality on the grounds that it does not provide any basis for the standard arguments in favour of semantic compositionality, which we discuss in §10. The claim is that, while these standard arguments allege that semantic compositionality is the *only* way to account for certain features of human, natural-language abilities, the definition in Definition (1) provides at best a very partial explanation.

6 Some Empirical Problems for Semantic Compositionality and some Suggested Solutions

The literature is rife with many different challenges to the adequacy of semantic compositionality to account for meaning in natural language. So many that there is no way to even seriously sample the landscape. In this section I merely consider a group of often-asserted

challenges to semantic compositionality. Their underlying commonality is that they report features of natural language that are allegedly impossible to be treated in a compositional manner. Although there is a certain further similarity amongst them (one might say that they are all manifestations of "the problem of context"), I do not pursue this, and treat them just as isolated topics. I furthermore do not report on the details of certain mathematical/logical sides of these treatments of the topics, but instead stay at an informal level of description and just give references to the more formal works. I leave out discussion of a number of other cases that have sparked much debate, such as the topic of noun-noun compounds (also known as conceptual combination in the psychology literature), issues about *any* and its interaction with negative polarity as well as its possible role in information friendly logic, and the general topic of possessives/genitives. Works that discuss many different examples, and give somewhat different takes on some of the examples considered here, can be found in Janssen (1997: pp. 437–447), Pelletier (1994b), Zimmermann (2012), and Dever (2006) (especially "the extended cut" version referenced below under Links to Digital Resources).

6.1 Context

There are two forms of "context" relevant to discussions of compositionality: one form appears to show cases where the meaning of some term X depends on or is determined or affected by some other linguistic item that X does not syntactically dominate. In such a case X's meaning would not be a function of the meaning only of its syntactic parts and the way they are put together. The second form is the sort already described in §5.1, where it seems that the meaning of some syntactic items (including whole sentences) depends on the non-linguistic context in which the items were uttered.

The present sub-section concentrates on the second sort of context. Different versions of the first type are featured in many of the remaining sub-sections of the present section.

One way to deal with the issue of context (of both types) is to make a very firm and sharp boundary between "literal meaning" and "conveyed meaning", saying that semantic theory is concerned only with literal meaning, while it is the subject matter of pragmatics to study all the myriad of ways that other meanings can be conveyed – such as by facial expression or kicks under the table, or .... One might call this attitude *strict semantic minimalism*. According to strict semantic minimalism, there is no problem of context for semantic compositionality, since contextual information is not a part of semantics – literal meaning can always be stated compositionally.

A somewhat different response is again to claim that there are two levels of meaning: one corresponding to "literal meaning" but the other corresponding to a "fixation of meaning" for the purposes of a specific discourse. Both levels are to be included in the meaning of an utterance (now in a specific discourse), and both should be given a compositional treatment.

In this view, certain terms are assigned a "localized meaning" and the various sentences of the discourse have themselves a localized import or meaning in that situation. It seems plausible to call this local meaning that is relative to a situation, the "situated meaning". This is *not* to include everything that might be included in the more general "conveyed meaning", but only a certain more-or-less agreed-upon set of fixed assignments to the various basic words (which is perhaps a part of the "common ground" of any linguistic interaction). The response then claims that *both* the literal meaning *and* the situated meaning are compositional: if one considers only the literal meanings of the atomic terms and the literal semantic effect of the various syntactic rules, then the resulting literal meaning of the sentences can be compositionally derived. On the other hand, if one considers only the situated meanings of the atomic terms and whatever might be a situated meaning of some syntactic rule (for instance, maybe sentence-final raising intonation (upspeak) is a syntactic rule that has some situated semantic effect), then one can compositionally derive the situated meaning of sentences. The alleged violations of compositionality are due to the mistaken attempt to use the literal meanings of the components to generate the situated meaning of the sentence.

Consider a sentence

(6) That prospect is very big.

It should be expected that, from the literal meanings of the various terms and the literal effect of that predication, compositionally deriving a situated meaning that some rookie basketball player is very tall will not be possible. Nor, it goes without saying, will it be possible to derive a different situated meaning that some business deal will reap huge rewards. But were one to *start* with the different possible situated meanings for *that prospect* and *big*, one could plausibly derive those situated meanings compositionally. The question then becomes whether it is possible to give situated meanings to the more basic syntactic items, without having to first derive a full sentential literal meaning. A general name for this possibility seems to be "interactionist semantics" (see, e.g., Recanati (2012: p. 178f)), where "simple representations [i.e., of lexical items] have their contents determined, in part, by the complex expressions in which they occur." Recanati holds that his notions of "modulation" and "free enrichment" can be treated this way.

Making that general idea take on a more formal flavour can be done by treating the situation as having two different aspects: what are called 'context of utterance' and 'circumstance of utterance'. Since a wide variety of theories can make use of the general framework, we see different theories putting different of these items into each of these categories. Kaplan (1989) makes contexts contain the *agent, time, position, (possible) world*, while circumstances are pairs of times and worlds. But there are other ways to carve up the things being used to characterize a speech situation.

In Kaplan's formulation there is a distinction between the level of *content*, which is obtained by fixing the utterance context (but not the circumstances), and *character*, which is the function that takes expressions *to* contents. It is worth mentioning that Kaplan's "Ban on Monsters" in effect says that contents behave compositionally (as do characters, naturally).[3] But Kaplan's is just one of many ways to conceptualize the general framework. (Bigelow 1975a,b) describes a view according to which linguistic items can enlist very many different sorts of things from the local context. This sort of enlistment he calls "quotation", although it is much broader than just what we ordinarily envisage by quotation. In this sort of picture, linguistic items "co-opt" items from the context and carry them into the semantic structure that is to be evaluated. Thus *him* could bring a particular person into the semantic representation; *nearby* might bring some distance that is presumed in the specific speech situation to count as nearby. Bigelow's picture is that although the "unadorned sentence" will fail the context-test for compositionality, once it is augmented with these co-opted items from the context, the resulting semantic representation *will* be semantically compositional. On the other hand, unlike the strategy taken by Kaplan and others, Bigelow does not think that there is some finite list of contextual features that can be identified univocally, but rather he sees this as a part of pragmatics: anything that can establish a mutual knowledge of who/shat the utterer wants the audience to infer that the uttered in mind. Although there is no general recipe, in particular situations, the utterer can count on the audience to know which referent or disambiguation (etc.) the speaker had in mind. So in a way, Bigelow uses a part of the strategy of the semantic minimalists, with a strict distinction between semantic and pragmatic information, but imports the pragmatic information into the semantic representation. (So this is also a type of "interactionist semantics", in the terminology of Recanati, that we discussed just above.) Westerståhl (2012) contains a characterization of a very wide variety of such conceptualizations (including such interesting cases as when the circumstances include "judges" who evaluate "statements of taste" in accord with relativist semantic intuitions), and gives a formal account of the various context-dependent forms of compositionality that are possible in them.

Westerståhl concludes his very nice examination of the interaction of context and compositionality with (pp. 218–219):

> The pervasive context dependence of natural languages, in all its forms, may seem to conflict with compositionality, or systematicity. Hopefully, the observations in this chapter can alleviate such worries... We have seen that compositionality and context dependence are not incompatible, indeed that contextual compositionality of content, far from being opposed to traditional compositionality of character or standing meaning, in fact *entails* it, and that ... versions of contextual compositionally relate to how semantic

> theories deal with things like context shift failure, unarticulated constituents, modulation, etc. ... In sum, the results here ... show in what forms and for what kinds of semantic values one may reasonably raise the issue of compositionality, when extralinguistic context is taken seriously.

Anyone who is concerned with the influence of context on semantic compositionality – regardless of whether it is for them a worry about their favourite compositional semantics or that they believe context will undermine semantic compositionality – needs to study Westerståhl (2012).

6.2 *If* and *Unless*

The case of the interaction of *if* and *unless* stems from Higginbotham (1986), where the discussion is about the peculiar interaction of the quantified subjects and the connective *unless* in sentences such as

(7) a. Every person at the party ate steak unless they ate lobster.
b. No person at the party ate steak unless they ate lobster.

Intuitively, (7-a) means that everyone at the party either ate steak or ate lobster, whereas (7-b) means that everyone at the party who ate steak also ate lobster. The semantic representations provided by Higginbotham (using his notation for generalized quantifiers) was

(8) a. <*every x*: Person(x)>(Ate-steak(x) $\vee$ Ate-lobster(x))
b. <*no x*: Person(x)>(Ate-steak(x) $\wedge\neg$ Ate-lobster(x))

Note that *unless* in both sentences in (7) does not dominate the quantifier phrases, so the interpretation of *unless* cannot compositionally employ information about which quantifier phrase is in force for its interpretation. Yet, as the formal representations in (8) show, when *unless* is in the scope of *every* it means "$\vee$" while when in the scope of *no* it means the truth-function "$\wedge\neg$". So its interpretation requires information from linguistic context that it does not dominate, and thus is a violation of compositionality.

Pelletier (1994a) challenged this argument by giving two different, general ways to return to semantic compositionality for any apparent counterexample of this general form. I think the main moral of the dispute is to see that there always are two such general approaches that might be taken in the face of many similar examples (and have been taken over and over by semanticists hoping to re-establish compositionality). The first method is what Pelletier called "the semantic solution", and it consisted in making the initial (lexical?) meaning of *unless* be "vague" or "underdetermined", or more formally,

be "some connective from among this set" (For the example at hand, the set was just two-membered: $\{\vee, \wedge\neg\}$, i.e., '(inclusive) or' and the truth-function 'and-not'). When this underdetermined connective is employed to connect two VPs, we have a lambda abstraction of this meaning. This very same undetermined meaning of the conjoined VPs would be generated (compositionally) no matter what the subject term was. And now it is the task of the final combining of the NP subject with the conjoined VP predicate to particularize the underspecified VP meaning correctly. A natural way to do this is to assign the *no* generalized quantifier a semantic feature of 'negative' and the *every* generalized quantifier a semantic feature of 'positive'. In the presence of a 'positive' NP that is combining with this underspecified VP, the connective is realized as $\vee$; in the presence of a 'negative' NP it is realized as $\wedge\neg$. This is completely compositional: no semantic rule to interpret a syntactic node depends on any information other than that of the node's immediately-dominated nodes.

The second solution is called "the syntactic solution" and amounts to varying the syntactic categories involved in the lexical items. The particular way it is done in Pelletier (1994a) is to make *unless* be a homonym, differentiated in the lexicon with syntactic features [$\pm$neg]. Correspondingly, the quantifier lexical terms, *no* and *every*, are assigned the [$+$neg] and [$-$neg] syntactic features respectively. And to the syntactic rule that checks (among other things) whether both the subject NP and predicate VP have the same [$\pm$pl] syntactic feature, we add that they have to have the same [$\pm$neg] syntactic features. Once again we will get the results in (8) that Higginbotham wants.

As I said, the moral seems to be that problems for semantic compositionality – at least of the type that looks to the wider linguistic context (if it is in the same sentence, anyway) for a solution – will admit of two types of solution: a semantic solution where the meaning of some item is said to be "vague" or "underspecified", only to become "specific" at some later stage by dint of combining with particular other meanings, and a syntactic solution that uses "ambiguity" by multiplying the number of terms so that each of these terms can bear a marker for just one of the possible ways to be combined syntactically.

6.3 Adjectives

The literature on adjectives and the many difficulties they produce for compositionality probably started when it was first noticed in the formal semantic literature that meanings of adjectives vary depending on what noun is being modified. In the earliest version of the problems, the "context" is linguistic and the adjectives are scalar: *big* when modifying *problem* seems to mean something different from (and maybe even incommensurate with) *big* when it modifies *house*. Later, but especially with the appearance of Lahav (1989), it was argued that (almost?) all adjectives show this property, and not just scalar ones.

Lahav produces a wide variety of examples featuring (mostly) *red* but also mentioning an impressive number of other adjectives (a small sample: *good, pretty, blunt, flat, square, slow, sad, strong*). The point of the examples was to show that the "applicability conditions" of these adjectives can vary with no regularities at all when applied to different nouns. Thus, when for example *square* is applied to *face* the applicability conditions involve the contours of the chin, cheeks and forehead as they appear from the front. But when applied to *screwdriver* it applies to a square end. The applicability of *sad* to *story* concerns tragic events whereas its applicability to *voice* describes intonation; both differ from the applicability conditions of *sad* as applied to *person*.

From the wide variety of examples where the applicability conditions differ so fundamentally, Lahav concludes that the variance of the semantic contribution of adjectives cannot be analyzed in terms of any general rule, unless it is "vacuously disjunctive".

Reimer (2002: p.190f) objects to Lahav's characterization

> Lahav is led to error because he fails to distinguish between L[iteral]-meaning and C[ontextual]-meaning. ... Lahav's mistake is in assuming that *content* must be uniform across contexts, if adjectives are to conform to compositionality (in any substantive sense). But this is a mistake. *L-meaning* (or conventional meaning) is what is uniform across contexts; content varies across contexts. This is just as it is with run-of-the-mill indexicals like 'I': L-meaning is constant. That what it means for the expression to be *univocal* (rather than ambiguous). In contrast, the C-meaning of an indexical expression varies with context. Such is the nature of an indexical (context sensitive) expression – whether it be an adjective, a pronoun or an incomplete definite description.

Sainsbury (2001) takes a somewhat different view, thinking that adjective-noun combinations should be viewed as "unspecific but determinate and unambiguous meanings".

> 'Italian F' is satisfied by a satisfier of F of or pertaining to Italy. This is an unspecific but definite, unambiguous and complete meaning. An Italian book is one of or pertaining to Italy, and a book may pertain to Italy by being on my pile of books to take there, or by being about Italy, or by being manufactured in Italy, or by being written in the dominant language of Italy. ... 'Feline care' is satisfied, in the same sense though not in the same way, by a vet tending a cat and by a cat tending the puppy to which she is acting as foster mother. Both vet and cat supply care of or pertaining to felines ...

Sainsbury's main point in this is the contention that theorists should distinguish between different 'readings' of a sentence from different ways in which it could be made true.

Just because (for example) *This is an Italian book* can be made true by this book's being about Italy, or written in Italian, or the book bought while visiting Italy, and so forth, it does not at all follow that this shows that there are somehow different 'readings' or 'meanings' involved with the sentence. And if these are all the same 'reading'/'meaning', then a compositional semantics needs only to generate a meaning for *Italian book* that is a function of the meanings of *Italian* and *book*.

A different answer to this sort of "effect of the wider linguistic context" issue has long been taken by formal semanticists. It is that adjective+noun combinations are not to be analyzed in the way given in elementary logic textbooks, as a conjunction of the adjective meaning with the noun meaning.[4] Instead, adjective-meanings are interpreted as functions that operate on the meanings of the nouns (including complex common noun phrases) they modify and the result is a meaning that picks out some subset of the interpretation of the noun (thus, a so-called "surjective" adjective). The meaning of *red* in *red*+NOUN is a function that selects out some subset of the interpretation of NOUN; thus, the meaning of *red apple* is the subset of apples that are red, the meaning of *red wine* is that subgroup of wines that are red, the meaning of *red hair* is ... etc. That is to say, we have *apples-that-are-red-in-the-apple-way* or *human-hair-that-is-red-in-the-human-hair-way* or *wine-that-is-red-in-the-wine-way* or *grapefruit-that-are-red-in-the-grapefruit-way*. In the symbolization of Parsons (1968, 1970), we have $Talented(Violinist)(x)$ rather than $Talented(x) \wedge Violinist(x)$. (If you had the latter, you would also know of anyone who was a talented violinist and a linguist, that she was a talented linguist!) Of course, this suggests that a part of what one needs to learn about *Talented* – in addition to the meaning, which is that it selects a subset of the noun it modifies – is just exactly *what* subset is picked out in all the different cases. But this does not imply that there is nothing in common in all these different cases, only that the details of picking out may differ among the cases. But surely that's as it should be! A child's learning that someone is a talented violinist and is an artist should *not* allow the child to know that the person is a talented artist. The child surely needs further, different information for that! (Parsons (1968, 1970) also uses this type of interpretation to accommodate adverbs as well as adjectives.)

This is indeed the type of solution most commonly proffered by advocates for semantic compositionality to the issue of adjectival modification of nouns. The literal or standing meaning of *red*, for instance, would be as in (9-a); and when this meaning is applied to the meaning of *apple* (which we represent as APPLE), we get (9-b), which yields (9-c) the property of being a red apple.

(9) a. $\lambda X \lambda x \{X(x) \wedge (\text{RED-FOR-AN-}X)(x)\}$
b. $\lambda X \lambda x \{X(x) \wedge (\text{RED-FOR-AN-}X)(x)\}(\text{APPLE})$
c. $\lambda x \{\text{APPLE}(x) \wedge (\text{RED-FOR-AN-APPLE})(x)\}$

But even in the reasonably simple case of scalar adjectives, like *big*, this solution doesn't always seem to work: Heim & Kratzer (1998: p.71) illustrate the problem with this story,

> Imagine we had first introduced a scenario populated with an army of monsters like King Kong. We might have said something like 'Jumbo doesn't have a chance; he's only a small elephant'; and this could have been true even if Jumbo were as large as or even larger than most other elephants.

Here the class that is relevant to the comparison class is not elephants but rather the class of monsters like King Kong. So the function here does not employ the class indicated by the noun, but rather a class that is determined by the context. Once this is pointed out, various other examples in the literature – such as the examples with *cut* in Searle (1980) or the examples of painted leaves in Travis (1997) – can be seen as making this same point.

But since this is now seen as reverting to the use of extra-linguistic or "common ground" contextual information, the general methodology described in (9) can be kept but with the idea that we will need to employ contextual factors of the sort described by Westerståhl (2004) and recounted above in §6.1. It seems then, that the wide range of semantic flexibility shown by such examples points not to a violation of semantic compositionality but rather to the need for a widening of the notion of semantic evaluation function to accommodate the "Kaplan-style" semantics suggested by Westerståhl.

6.4 Ambiguity

Generally recognized in the linguistics-philosophy literature are two types of ambiguity: lexical and structural. One might add to this another, which is often lumped together with lexical ambiguity – homographic or homonymic ambiguity. Because of the graphical distinction, no one would accuse

(10) a. John adjusted the sink in order to smooth out the sound.
b. John adjusted the sync in order to smooth out the sound.

of being ambiguous (at least not on account of *sink/sync*), although a phonetic representation might be thought to be. But when the relevant items are spelled identically, there is a temptation to see this as a sort of lexical ambiguity, as in

(11) The bark was annoying and even caused pain to some of us.

where *bark* might be the tough, woody exterior of a tree or bush, or equally, the sound made by a dog. But this is not a lexical ambiguity in any straightforward sense, because these are two different words, spelled and pronounced identically. And similarly with the

phonetic representation of (10) – despite the identity of pronunciation, it is not a lexical ambiguity because these are two different words.

A real lexical ambiguity is *one* polysemous word, for example

(12) The puff surprised John,

which may describe John's surprise at an abrupt emission of air or at an exaggerated praise of a book or . . . (each of which describes a different meaning of the one word, *puff*, in the Oxford English Dictionary). And furthermore, *surprised* in (12) refers to the feeling of astonishment felt by John by something unexpected. It does not make it a different meaning that this feeling was occasioned differently (by feel vs. sight, for example), nor that the level of astonishment was different, nor that John might do different things as a result of the surprise. There are many, many different ways that one and the same meaning might be instantiated; many, many different ways that one and the same meaning might have different "applicability conditions". By themselves, these differences do not make for a difference of meaning – something that should be kept in mind by those thinking to overturn compositionality on the basis of "different applicability conditions".

Since the discussion is about *semantic* compositionality, and since one aspect of this area concerns ambiguity, it is surprising that not much has been said in this literature about the difference between polysemy and homography/homonymy and whether or not this has an effect on semantic compositionality. It seems to me that it is not clear at all how the various movements in linguistics intend to formally incorporate the meanings of lexical items. If it is just that the formal structure (a tree, or whatever) will have a *string* from the lexicon be inserted appropriately (as a leaf node of a tree, or whatever), then one will see little difference between homography and polysemy. But given an antecedent preference for distinguishing distinct *words* – or rather, distinct lexical items – then one will want to make that distinction. However, it then becomes a challenge to see how to characterize the insertion of one meaning from insertion of a different meaning of the same word in a sentence. Consider

(13) He hit a pitch over the wall

which might be a claim made in a baseball or cricket game. In these cases it would use the same meaning. But it might also be a claim made in a golf game, where it employs a different meaning. It is quite difficult to see how a process-oriented model for choosing the correct meaning would work. And if one does not even distinguish homography from polysemy, it is equally difficult to see how the different sentences represented within (10) or within (11) would be handled, as they surely must be.

The second established type of ambiguity, structural ambiguity, occurs when the same lexical items are given different structural descriptions by the syntax of the language, as in

(14) a. old men and women
b. the shooting of gang members
c. flying airplanes

(Are the women supposed to be old? Or just the men? Were the gang members doing the shooting? Or being shot? Are we talking about the activity of piloting an airplane? Or about airplanes that are airborne?)

Both types of ambiguity are thought to be easily handled in a compositional semantics. Lexical ambiguities are to be dealt with right at the start, by assigning the appropriate meaning to the lexical item and using these two different meanings to generate two meanings of the larger syntactic item. Similarly with structural ambiguity: since different syntactic rules come with distinct semantic functions, it is clear that structural ambiguity can yield a pair (or more) different meanings without any violation of compositionality. If there is some (non-semantic) reason to know that one or another of these meanings is inappropriate for the case at hand, that is not an issue for compositionality but rather for some other component.

However, there are a number of cases that have been thought to be problematic. One clear case is that of quantifier scope ambiguity. Everyone seems to agree that sentences such as

(15) Someone is hit by a car every day

are ambiguous in that there are different understandings of where the quantifiers scope with respect to one another. Different possibilities can be paraphrased

(16) a. There is a person and a car such that every day the person is hit by that car
b. There is a person such that every day that person is hit by some car
c. There is a car and a person such that every daythe person is hit by that car
d. There is a car such that every day that car hits somebody
e. Every day there is a person who is hit by a car that day
f. Every day there is a car that hits some person that day

Not all these understandings are logically distinct: (16-a) and (16-c) are equivalent; (16-e) and (16-f) are equivalent, and various ones entail others. It is some sort of "world knowledge" that makes us gravitate to one of (16-e) or (16-f). So (15) is ambiguous between the various non-equivalent understandings, and thus compositionality requires some sort of difference in its lexical items or its structure, so as to give a foothold for the ambiguity.

Most of the semantic theories that have been added to traditional Chomskean theory have advocated a syntactic level of LF, where the quantifiers will have been placed appropriately. And therefore in the present sort of case, there needs to be some options open in the syntax

to produce different placements for the quantifier phrases so as to have different LFs for (15).

The theory of direct compositionality (see §8.3) argue strongly that this sort of syntactic manipulation robs compositionality of its "bite", and that it is an "unprincipled" way of envisioning semantic theory. Be that as it may, many others have claimed that there is simply *no* syntactic ambiguity in (15), and that also raises problems for compositionality, since the straightforward semantic representation (which would follow the left-to-right occurrences of the quantified NPs in (15)) would yield the meaning expressed in (16-a)... which no one thinks is right.

Long ago, the notion of "Cooper storage" was brought in to deal with this problem (Cooper 1983). The underlying idea is that quantified NPs are held "in a storage facility" while the remainder of the sentence is syntactically generated and semantically evaluated (to the extent that it can be evaluated with only placeholders rather than explicit NPs). Then the NPs are parcelled out to the various placeholders. In a "pure" version of this, any of the quantified NPs could be the replacement for any of the placeholders, although most versions also had some other constraints about matching features like plurality and gender. In our example, this would generate all six of the meanings in (16), thus showing an underlying six-way ambiguity in (15). To many theorists it was not completely clear that Cooper storage really was in accord to compositional principles, but it was widely adopted nonetheless by compositionalists.

Pelletier (1994b, 2000) argued that such sentences "really" have only one syntactic structure, that Cooper storage was not "really" compositional, and that therefore compositionality fails for them. Furthermore, there were other sentences that displayed what he called "serious ambiguity" (ambiguity not traceable to any lexical or structural feature), such as

(17) a. John wondered when Mary said she would leave
b. When Alice went to school, she rode her bicycle
c. The Canadian family used less water last year than the preceding year

In response, (Pagin & Westerståhl 2010b: pp. 21-22) remark that there are other options, such as having *unspecified meanings* of some kind, and treat the phenomenon under discussion as a kind of contextual issue. (And employ some of the suggestions from Westerståhl (2012) to carry this out.) Alternatively, we might wish to keep the specificity of sentences, but make the meaning assignment be "relational", that is, make sentences be possibly related to a *set* of meanings. Although this is not compositionality in the sense of Definition (1), perhaps this is the way to go. (One would want to start anew in the justification of semantic compositionality, making sure that it still is distinct from (e.g.)

semantic holism, that it is still "aesthetically pleasing", and that it continues to justify the arguments for compositionality discussed in §10.)

6.5 Idioms

Given that the very definition of an idiom is as a phrase whose meaning is not predictable from the literal meaning of its parts, it seems strange that some semanticists should try nonetheless to give a compositional semantic for idioms. Yet some theorists have sought to do so. Here I just describe one line of thought.

It seems to be a standard "first try" to assimilate idioms to the category of lexical atoms. In such a case idioms like *kick the bucket* and *trip the light fantastic* would each be a single lexical item and assigned the meanings DIE and DANCE, respectively. Of course, there are also literal uses of *kick the bucket* (and possibly of *trip the light fantastic*), so there would be ambiguity in the sentences that contain the phrase. But this type of ambiguity would not be telling against semantic compositionality any more than lexical and structural ambiguity is.

A problem arises, however, when it is noticed that many idioms seem to participate in syntactic rules while nonetheless maintaining their idiomatic meanings. Standard examples are

(18) a. Strings were pulled by her father to get Anita her job
b. We will leave no legal stone unturned

and very many more are explored in Nunberg et al. (1994). The point is that such idioms can't be treated as "atomic elements" of a lexicon, since they *do* participate in these more complex structures, following standard syntactic rules that apply to their parts. In fact, even the most "semantically opaque" of idioms seem to participate in at least some syntactic rules. *Kick the bucket* and *trip the light fantastic*, among the most "atomic" of idioms, require verb agreement with the subject and tense/aspectual agreement with other parts of the sentence in which it occurs.

(19) a. John kicked the bucket yesterday
b. Everyone eventually kicks the bucket
c. Aunt Martha will surely kick the bucket before next week

(20) a. John and Mary were tripping the light fantastic yesterday
b. Tonight, in my fantasies, George and I will trip the light fantastic

This suggests that the lexical atom approach is not really viable. However, it doesn't rule out the idea that the elements of the idiom…the lexical parts of the idioms…can't

have separate idiomatic meanings in the lexicon. *Strings* and *pull* could have two different meanings (or be two different, homographic, words), and participate in the same syntactic rules/descriptions, but the differing starting meanings could give rise to different meanings of the sentences in which they appear. And both sentential meanings would be generated compositionally. (But it doesn't follow, as Nunberg et al. (1994) are at pains to point out, that a person would necessarily know both the literal and idiomatic meanings of the words. And thus it doesn't follow that a person would understand the idiomatic meaning even were the non-idiomatic meanings of the lexical parts known, and even were the these meanings compositionally composed to generate a non-idiomatic meaning of the sentence, which the person knows.)

The approach in Westerståhl (1999) is to ask of a language that has a compositional semantics, whether a new idiom (using the lexical items of the language) can be added to the language and still preserve compositionality. His work is, as he emphasizes, general and abstract, where the idea is to be able to show that a wide variety of different conceptions of language and a wide variety of different conceptions of how idioms behave can be all given some compositional treatment or other. Different specific grammars will have more tools at their disposal than is presumed in the abstract treatment, and so they may offer specific refinements of the abstract guarantee of the existence of a compositional extension that includes idioms.

Here is what Westerståhl (1999: p. 175) says about his method and results.

> One suggestion here is that *kick the bucket* is somehow an atomic expression, in contrast with *pull strings*. A different idea is that the former idiom stands for a one-place predicate, the latter for a two-place one, so it is natural that passivization and anaphoric reference to the object position apply just in the latter case. This kind of semantic explanation may (on some accounts) be applied regardless of whether an idiom is seen as having structure or not. If it does have structure, there are still various ways to think about it. One natural idea is that (some) idioms have syntactic but not semantic structure, as it were. An alternative is that they have both. Then (some) idioms are composed of parts (some of) which themselves have idiomatic meanings. All of these suggestions will be considered in this paper.

Speaking generally, Westerståhl's analysis allows for some idioms to be treated as "atomic" and unanalyzed, but for others to appear in the same syntactic constructions as the corresponding non-idiomatic expressions do, while computing their meanings using the familiar semantic operations. The "trick" for such an analysis is to assign the correct idiomatic meaning to the subparts of the idiomatic expression. Westerståhl (1999) thus follows the natural understanding of Nunberg et al. (1994), where (certain) idioms are

viewed as atomic lexical items and given their idiomatic meaning in the lexicon while others are syntactically complex but their lexical components are "duplicated", with one being marked as idiomatic. Westerståhl remarks on the HPSG implementation of the Nunberg et al. (1994) idea as it occurs in Sag & Warsaw (1999: Ch. 11), saying that since there are many particular aspects of HPSG that are not presumed in his (Westerstaåhl's) abstract discussion, there are a lot of ways that the HPSG framework can work to avoid the possible "overgeneration" of idiomatic meanings that the abstract framework would allow. Westerståhl (1999) ends with this claim: "I believe there is one conclusion that may be safely drawn from the formal work in this paper: With respect to the principle of compositionality, idioms do not pose a particular problem."

6.6 Quotation

There are three commonly-identified types of quotation, and the details of how to distinguish them can be quite theory-bound. (A survey of the various types of quotation and differing analyses of the different types is given in Cappelen & Lepore (2012).) However, without trying to be too precise we can say: *pure quotation* occurs when a linguistic item is being "directly" referred to, *indirect quotation* occurs when one is reporting (in one way or another) without exactly naming the item being referred to, and *mixed quotation* occurs when there is a mixture of the two. In orthography pure quotation conventionally employs quotation marks, indirect quotation makes use of *that*-constructions (and others), while mixed quotation usually employs quotation marks within a that-construction. (21-a) illustrates direct quotation, while (21-b) and (21-c) illustrate indirect and mixed quotation respectively.

(21) a. 'Quotation' has nine letters
b. George thought that indirect quotation would make for a good thesis topic
c. Alice told George that indirect quotation would be "a hard sell" to the committee

This section is concerned with pure quotation, even though the main work we discuss offers a framework within which all three might find a home (Pagin & Westerståhl 2010). Now, in a way that is similar to the reaction one might have upon hearing about semanticists' wish to give a compositional treatment of idioms, so too might it strike the reader as strange that some semanticists should attempt to compositionally describe the meanings of the pure quotation of linguistic items – after all, isn't this a case where the *meaning* of the quoted material is *ir*relevant to the designation of the quotation as a whole?

Well, it is obvious that if (a) the phrase being quoted is a syntactic constituent of the quoting phrase[5], and if (b) at least some quoted phrases are semantically equivalent to distinct other phrases, then the language cannot be semantically compositional. For, this

would be a direct violation of substitutional compositionality (from §8.1); and from there, with a principle that syntactic parts of meaningful parts are also meaningful, we get a violation of Definition (1), Functional Compositionality.

So obviously, any theory that wants to give a semantically compositional theory of pure quotation will have to deny one or the other of (a) and (b). Some theories re-analyze quoting expressions so that the quoted phrase is *not* a constituent of the quoting expression, thus giving up on (a). Theories that analyze quoting expressions as *names* (of the quoted expression) are one example, since names are an unanalyzed unit. Theories that analyze the quoting expression as *descriptions* (of how to form the quoted phrase from letters/words by concatenation) are another. A third example are those theories that treat quoting expressions as *demonstratives* (pointing to a token of the type being demonstrated).[6]

A different sort of response to the (a) vs. (b) approach is to employ what Pagin & Westerståhl (2010) call a *use theory* of quotation, according to which takes quotations to be indicators that we are to expect a different *use* of ordinary words. A central idea in use theories is that certain pieces of language (words, phrases, punctuation) indicate that other pieces of language should be "understood as being used in a non-standard way." And that these indicators might not be a part of the piece to be used differently. So this is a violation of semantic compositionality, at least if these "different ways of using X" amount to different meanings of of X than if they were not given such an indication.

An example that Pagin & Westerståhl (2010) use to give an informal explanation of their idea (and also to indicate that the method being advocated can be employed in other types of quotation than just the pure) is an interpretation of Frege: in a "belief context" the interpretation (*Bedeutung*) of the embedded words should be the sense that these words have when not embedded inside such a construction. To take this seriously, one has to isolate the fact that it is a belief context when evaluating the embedded words/phrases. Their idea is that this would be done by noting that in some *different* area of the sentence than the one being evaluated, there is the word *believes*. Having noted that, we then proceed to evaluate the embedded context in terms of the *senses* of its parts, rather than the more usual *Bedeutungen* of its parts.

One can see a certain similarity in this idea to the use of non-linguistic context as a way of determining the semantic values of various linguistic items being evaluated, and indeed, Pagin & Westerståhl explicitly draw this analogy. And although the use of extra-linguistic context is usually seen (by compositionalists, anyway) as not a challenge to semantic compositionality, the proposed use of areas of the linguistic item other than what is under evaluation strikes many theorists as dubiously compositional. (Although one might recall the two suggestions in §6.2 – one syntactic and one semantic – as a way to implement the idea in a way that seems quite naturally compositional.) Pagin & Westerståhl choose to call

their method one of "general compositionality", and show what the formal background for it would look like.

In pure quotation, their idea is that the quotation marks would function like the *believes* in the Frege-inspired example. But rather than saying that the semantic theory should use *senses* of the quoted terms, its interpretation should be the *name* of the quoted phrase. To carry this out, the initial language, assumed to be semantically compositional in the standard sense, is enriched by adding quote marks and a relevant syntactic rule that will add quote marks around any expression of the language and form a name in the language. And there is a corresponding semantic interpretation rule that the name thus formed is interpreted as the string formed. This will give a "g[eneral]- compositional" theory of pure quotation that allows for the quoted material to be a syntactic component of the quoting phrase and for the quoted material to be capable of semantic equivalency with other, distinct phrases (when not quoted).

The case of pure quotation has some unusual aspects, however. One is that one can cite pretty much *any* language-oriented item, not just pieces of the language being analyzed. Pagin & Westerståhl (2010) mention

(22) a. 'farfalla' is Italian for butterfly
b. 'str jd e' is not a sentence
c. 'א' is a Hebrew letter

Note that none of the quoted expressions in (22) is a term of English (which is the language of sentences in (22)). So, some further extension needs to be done to accommodate this, and (Pagin & Westerståhl 2010: §7) mentions a few ways of doing this, but remark that their abstract approach leaves open many different ways to incorporate this issue. They also say

> As we have emphasized, the extended semantics . . . is *not compositional*: one cannot substitute synonymous terms in the scope [of the quotation marks] and expect meaning to be preserved. However, [their method] is g-compositional . . . [our newly-defined meaning function] only switches [away from the normal meaning] when something is quoted.

They end with the claim that "even for the simple case of pure quotation, this kind of compositionality is all one can hope for. . . . [And] linguistic context dependence as outlined here can be exploited for other constructions, notably various kinds of intensional contexts." This looks to be a confirmation of one's initial reaction that *of course* quotations are not compositional!

Well, one way to take this last claim is that semantic compositionality cannot deal with any of the types of quotation. On the other hand, semantic compositionality also can't

deal with non-linguistic context in a literal manner either, but rather it needs to bring in an adjunct to evaluating semantic content, as discussed above in §6.1. Most theorists don't find this latter extension to tell against semantic compositionality; is it the same case here?

I don't think so. In the case of quotation, the claim that g-compositionality is "all one can hope for" is said with the background of Pagin & Westerståhl's restrictions (a) that the quoted phrase is a syntactic part of the quoting phrase, and (b) that some syntactically distinct phrases that can be quoted are semantically equivalent to distinct other quotable phrases. But these restrictions have been denied. Any of the name or description theories mentioned in Footnote (6) above violate (a), and some of these, as well as other theories, violate (b) as well. So, there are other options that are available, if one is not wedded to conditions (a) and (b). In addition, Pagin & Westerståhl mention Potts (2007) as not being compositional (but rather is "recursive"), and Shan (2010) as violating both (a) and "the natural requirement" that two quoting expressions can't be semantically equivalent if the quoted expressions are distinct. One might think that one or another of these theories is preferable; that is, one might think that violating one or both of (a) and (b) is not as aesthetically displeasing as adopting g-compositionality. One might note that similarity between Pott's recursive semantics and the "general recursion" advocated in the Atomistic, but non-compositional, Semantic Groundedness theory expressed in §8.4.

7 Ontology: Parts and Wholes

It may seem strange to some that a discussion of semantic compositionality in linguistics should include a portion devoted to ontology. But I think that it will facilitate discussion of some conceptions of compositionality (and their opposites), as well as giving an overview as to what "really lies behind" adherence to one or another of these positions.

I start by discussing Atomism, and follow that with a discussion of Wholism.

7.1 Atomism

"All-in Atomism" is the (ontological) doctrine that there are smallest, indivisible items of all reality. It is not so clear that there are any current holders of All-in Atomism these days, but in any case the sort of Atomism we shall find interesting is where the notion is relativized to *a field of inquiry*. Thus, one might find biology to be atomistic if it holds that there are smallest kinds of living things (cells, or whatever) and that all other kinds of living things are somehow and in some way composed of them. A semantic theory might be atomistic if it holds that there is a set of minimal primitive meanings and that all complex meanings are some type of combination of these. In this conception of ontological Atomism, we have wholes of some field of inquiry seen as composed of parts (also from

that field) and these parts composed of subparts (from that field), until one reaches the atoms (of that field). The mantra of Ontological Atomism is

Definition 2. (Atomism Slogan) *A whole is a system consisting only of the materials in its parts, possibly rearranged. And there are smallest parts of any such system.*

It should be noted that this Ontological Atomism admits as parts of a whole only those items that are its parts (and their parts, etc.), possibly rearranged.

Definition 3. (Ontological Compositionality) *A whole is built up only from materials in the parts, possibly re-arranged or re-configured.*

This "Building Block" theory differs in its presuppositions from those of semantic compositionality (e.g., as given in Definition (1)), which instead presumes that there is an already-given, well-defined system of the sort that Ontological Compositionality has characterized – namely the syntactic structure of language. In this syntactic system, the atoms are lexical items and the wholes that are made from them are syntactically-defined larger and larger linguistic units composed from the smaller units. But functional compositionality instead asks us to focus on some *property* of this previously-given system or in some other things that are somehow *related* to this system[7] – for instance, in the *meanings* of the items. In this setting, the issue becomes whether there is a way to define this property or things – the meanings – of *all* complex items solely in terms of the meanings of the syntactically-given parts of the complex item and their syntactic method of combination.

An important difference between these two notions of compositionality – ontological and semantic – concerns whether some "whole" can contain things not in the parts. According to the ontological building-block view, no; but according to the semantic version, yes. For, the first notion allows the whole to contain only what is in the parts, possibly re-arranged or re-configured in some manner. But the second allows the object or property that is associated with a whole (in the linguistic case: the meaning of a complex whole) to be a *function* of the things associated with the parts (in the linguistic case: a function of the meanings of the syntactic parts and syntactic mode of combination). There is nothing to stop such a function from introducing new material into the thing associated with the whole — that is, the function can make the meaning of a whole contain many new and radically different things than are contained in the meanings of the parts. According to functional compositionally, all that is required is that this be a function — which, as we saw earlier, merely means that it must introduce the same material and generate the same result every time it is faced with the same parts and manner of combination.

Atomism in general is allowed to employ a functional notion of compositionality; that is, ontological compositionality is not the only option open to Atomism. In §8.4 we discuss one such different version; another is described in §9.2. The point of the current

section is merely to emphasize that the Building-Block Ontological Compositionality of Definition (3) is not the version of compositionality that is employed by semanticists. The objections to semantic compositionality of the form "such-and-so syntactically complex whole contains information that is not in any of its simpler parts" – which one can find in many writers, both within linguistics and also in cognitive psychology/psycholinguistics – is just beside the point as a comment on Definition (1).

7.2 Wholism and Holism

Wholism (and some accounts of holism) form the opposite ontological stance to Atomism. Wholism is the view that in the area under discussion, there are some items whose makeup at least partially consists in things that are not in its parts. (Recall Footnote (2) in which the OED's account of the origin of the terms 'wholism' and 'holism' was cited. It is not so clear that Smuts and Boodin really differed in their intended targets for the different terms: the notions of emergence and gestaltism seem to be a part of both doctrines. Some modern versions of holism, however, are rather different from wholism, and we will look at this shortly.)

Although usage is not fixed in the literature, I will here use the terms 'wholism' and 'holism' to mark a difference, as follows. The fundamental idea behind wholism is to give some sort of priority to a "whole" in preference to the "parts" of that whole, and of course this plays out differently in different areas where the term is used. One basic difference is between 'wholism' (the view that there are objects ("wholes") that need to be referenced as primitives in some realm of inquiry) and 'holism' (the view that properties of the individuals under consideration require reference to properties of other individuals before the original properties can be fully defined or explained) is parallel to that of ontological vs. functional compositionality within Atomism. (There are also attempts to justify wholism in some area on the grounds that wholes contain properties that none of the whole's parts have. But we've already seen that the similar argument in the Atomistic realm is invalid. So we don't address it here.)

The first sort, wholism, is:

Definition 4. (Wholism) *Wholism is the view that there are Primitive Wholes – complex items that cannot be defined in terms of their parts.*

Sometimes wholists – like Boodin – bring forward considerations of gestalt theory and emergent properties to bolster the claim that there are primitive wholes. The square that is formed (or *appears* to be formed) when very heavily-blackened cut-out corners are drawn is one such gestalt figure that is common in the literature. "Surely," it is claimed,

"the square is *not at all* to be analyzed as those four heavily-blackened cut-out corners." But equally surely, even the ontological atomist never claimed any such thing. Rather, the square requires that the corners be placed *exactly so* in relation to one another to make the square appear. The square could be claimed by atomists to be a *function* of the heavy corners plus their method of arrangement or combination. This seems to be a case where the wholist ignores the atomist's "and the way the parts are combined."

Similar remarks can be made about emergent properties, that is, properties that the wholes have that are not possessed by any of the parts. For example, hydrogen and oxygen combining to give off energy and to produce water. The claim of wholists is that such properties are not "sums" of the properties of the initial hydrogen and oxygen. And since they are not "sums" of the properties of the initial parts, it has seemed to some that this calls for a new whole that possesses this property. But we can see that such considerations do not really play a role here. An atomist would attribute the property to the *juxtaposition* of the hydrogen and oxygen (plus whatever else is needed to initiate the reaction).

The dispute between Ontological Atomism and Wholism is fundamentally ontological: Are there non-individual items in the world? Of course, such a dispute turns crucially on what is meant by '(non-)individual', and so it is rather surprising to find that many of the disputants agree on what is meant, or at least, do not seem to challenge each other about this. Here are some of the relevant principles that wholists believe:

(23) a. The whole is more than the sum of its parts
b. A primitive whole determines the nature of its parts
c. The parts cannot be understood when considered in isolation from the primitive whole of which it is a part
d. The parts of a primitive whole are dynamically interrelated and interdependent
e. Analysis as a methodology fails in the case of primitive wholes

And conversely, atomists deny these five claims. As I remarked above, the Ontological Atomism vs. Wholist dispute is rather outside the realm of semantic compositionality, even when taken broadly. So I shall not follow up on this side of the dispute, but the interested reader can check out Pelletier (2012) for a more in-depth discussion.

However, it is important to appreciate how wholism impacts any discussion of semantic compositionality. For, a number of the anti-compositional positions *start* with a claim that there are wholes which have a different status from "the sum of their parts" – as when it is claimed that a sentence's meaning is a "unified whole" that is different from a "list of the meanings of the sentence's parts", or again, that the meaning of a sentence depends on the larger context or situation or "form of life" in which it is to be found. Without any further claims about these concepts – context/form of life/etc., there is nothing that is opposed to an atomistic and semantically compositional account of such phenomena.

The second sort, holism, differs from wholism:

Definition 5. (Holism) *(Property) holism is the view that some* properties *can only be attributed to entities that are not individuals.*

Property wholism is asserted in very many fields of enquiry, especially in the social sciences and humanities. Businesses and corporations, for instance, are claimed to have duties and obligations (etc.) that are not duties or obligations of any individual within the business. And they can similarly have legal rights and legal constraints that do not devolve to any individual. Nations might have properties that are not properties of any members or collections of members of the nation. A nation can declare war, or decree that their borders be closed; but no individual in the nation can do so, not even the leader.

In arguing against Functional Compositionality, this sort of holism (which is not the only sort – a topic to which we will return in §9.1.3) seems to have made the mistake that we surveyed earlier in discussing Ontological Compositionality: the claim compositionality presupposes that a whole can't have a property unless it is "already present" in the parts. This claim ignores the way that *functional* compositionality works. As we have seen, a function can introduce new and different features to the complex item, so long as it is done functionally – always generate the same result when given identical inputs.

There are also other types of Atomism than the one that is presumed in the Definition (1) of semantic compositionality. We mention one such version in §8.4. Other one seem to be in the background of some solutions that we surveyed in discussing linguistic problems for compositionality – for instance in the discussions above about context (§6.1), idioms (6.5), quotations (6.6), and perhaps others.

8 Other Notions of Semantic Compositionality

There are various notions in the literature which are more-or-less similar to functional compositionality as defined in Definition (1). Often this means either restricting even further the allowable cases that count as compositional (as when we discussed in §5.2 various suggestions about how to restrict what sorts of functions were allowed); but at other times allowing more types of languages to be semantically compositional.

8.1 Substitutional Compositionality

One that is very similar – indeed, it is often seen as an alternative definition of semantic compositionality – is *substitutional compositionality*: Intuitively, the substitutional version of semantic compositionality says that if the semantic values of two syntactic items are

identical – i.e., the two syntactic items mean the same (or, "are synonymous"), according to the semantics – then the semantic value of any item containing one of the terms will be identical to the same item when that term is replaced by the other of the synonymous terms. To state this, we use $\mu(t)$ to stand for "the meaning of (syntactic item) t", and use $X[t]$ to stand for the notion "t is a (syntactic) part of X". Note that this latter notion allows for t to not only be an *immediate constituent* of X, but to be a more "deeply embedded" subpart. (In principle, this gives rise to a number of different versions of this substitutional version, differing on "how deeply embedded" a subpart is allowed to be. We follow up this suggestion in the next subsection.)

Definition 6. (Substitutional Compositionality) *if* $\mu(t_1) = \mu(t_2)$*, then* $\mu(X[t_1]) = \mu(X[t_2])$

Now, this is not *quite* equivalent to functional compositionality. Functional compositionality makes a presumption that needn't hold in the case of substitutional compositionality, namely that all the syntactic parts of any syntactically complex item are themselves meaningful. (For, consider the statement of functional compositionality: "The meaning of a complex item is a function of the meanings of its parts". So, this presumes that all the syntactic parts of any complex item are meaningful.) It is also presumed that if two syntactic items have the same meaning, then they can appear in the same syntactic positions. Hence there is a sense in which substitutional compositionality is a more general notion than functional compositionality. (However, both these presumptions are usually thought to hold almost by definition.)

It should be noted that as it is stated in Definition (1), the function need not restrict itself to *immediate* constituents of the term whose meaning is being defined. And in return, substitutional compositionality has allowed for substitution of items that are *not (syntactically) immediate constituents* of the item whose meaning is being determined. However, if we assume that all syntactically-well-formed items have a meaning, then if we restrict the substitutions in the definition of substitutional compositionality to just the *immediate* (syntactic) constituents, this is obviously implied by functional compositionality. It can then also be shown, by an induction on the depth of (syntactic) embedding of the terms whose meanings are being used for substitution, that this result will also hold for arbitrary embeddings of the syntactic subparts. So, given the presumption of meaningfulness of subparts, functional compositionality implies substitutional compositionality. To show that the reverse implication holds, we would need to show that, given we have substitutional compositionality with its notion of meaning-sameness guaranteeing substitutivity, we can define the relevant functions called for by functional compositionality. Although this can be shown, it is a rather more difficult construction which we will not consider here.

8.2 Strong and Weak Compositionality

There are no doubt many different thoughts about what should be called Strong – as opposed to Weak – Compositionality. Further complicating the landscape here is the fact that there are conflicting intuitions about what 'strong' vs. 'weak' might mean in this context. As I'm using the terms, 'being stronger' means that it allows *fewer* theories to be compositional, whereas being weaker means that *more* theories will be classified as compositional. In discussing the various strong/weak distinctions applied to compositionality, I focus here on two related notions.

8.2.1 Larson & Segal

In (Larson & Segal 1995: pp. 78ff) Strong Compositionality is defined thus:

Definition 7. (L-S. Strong Compositionality) *A system of semantic rules is Strongly Compositional if and only if all its semantic rules are strictly local and purely interpretive.*

The two conditions mentioned are explained:

> These rules not only interpret a syntactic node . . . in terms of its constituents; they interpret it in terms of its *immediate subconstituents*. . . . They never look down any deeper in the [syntactic] tree. Furthermore, these rules interpret structure given by the syntax, and they interpret *only structure given by the syntax.* They never introduce structure of their own.[8]

This is the notion of semantic compositionality that most formal semanticists and philosophers have in mind: the meaning at each node of a syntactic tree is given as a function of the meanings of that node's immediate constituents, together with information about the method by which that node was constructed. In contrast, one might think of allowing the function to not only employ the information about the immediate constituents, but also of all the subnodes of these immediate constituents, and their subnodes, etc., down to the atomic (lexical?) parts. This would be

Definition 8. (L-S. Weak Compositionality) *A system of semantic rules is Weakly Compositional if all its semantic rules are weakly local. (And they are purely interpretive?)*

A semantic rule is weakly local when it can access information about any of the nodes dominated by the one which it is interpreting. Again, a meta-theory that claims that all semantic theories must be Strongly Compositional is stronger than a meta-theory that allows for Weakly Compositional semantic theories because it disallows semantic theories that a Weakly Compositional meta-theory would allow. Or to put it the other way: the class of

semantic theories that a Weakly Compositional meta-theory allows for properly includes the semantic theories that a Strongly Compositional meta-theory allows for.

8.2.2 Levels of Semantic Compositionality

There is an obvious generalization of the weak-strong compositionality distinction:

Definition 9. (Strongly Compositional to Degree *n*) *A system of semantic rules is compositional to degree n if it can access information about any of the syntactic nodes that it dominates that are at most embedded to level n beneath that syntactic node.*

In this characterization, the L-S Strong Compositionality is the meta-theory of Strongly Compositional-to-degree-1. Their Weak Compositionality can't be easily characterized, unless we allow for a symbol like '∞' to be used in 'Strongly Compositional to degree ∞' to mean "as deeply embedded as possible, down to the atomic (lexical?) elements of any syntactic object."

The intuitive idea of "levels of compositionality" is pretty clear, even if it is difficult to see why any semantic meta-theory would choose a level other than level 1 or level ∞. Level 1 allows the meaning of a node to be a function of the meanings only of its immediate constituents, and their manner of syntactic combination. Level 2 allows allows the meaning of a node to be a function of the meanings only of its immediate constituents and their syntactic manner of combination and also of the meanings of the constituents of these constituents and *their* manner of syntactic combination. And so on. It can be seen that with an increase in the degree *n*, more and more information is available to use in a function that is to determine a node's meaning, thus allowing more theories to be called compositional.

As with the previous examples, allowing for the semantic rules to have more information available to them by allowing information from more and more deeply dominated nodes to be used, allows for more semantic theories to be counted as compositional. And hence is a weaker meta-theory. That strong compositionality asymmetrically implies weak compositionality is shown (by means of an artificial example) in Zimmermann (2012: p. 83f).

8.3 Direct Compositionality

Direct Compositionality (Jacobson 2002, 2012, Barker & Jacobson 2007) is not really a different view of semantic compositionality *per se*, but rather a viewpoint on how a grammar as a whole (crucially including both the syntax and semantics) should best be organized. The main opponent of this view is the attitude that the syntax should be charged with developing an ineliminable "level of interpretation" and that the semantic component

should then be applied to this level.[9] The more targeted version of this opposition theory is that associated with a level of "Logical Form", LF. In this more-common conception of the relation between syntax and semantics (more-common at least in the version of linguistic theory that has its roots in Chomskean syntactic theory), the syntactic component constructs and manipulates structures in such a way that there is ultimately a structure of LF. A main critique of such a conception, and one emphasized in the Direct Compositionality literature, is that it allows the syntax to "arbitrarily" develop the LF in such a way that it bears only a very distant relationship to the surface string that is to be interpreted.

Direct Compositionality then is the view that the semantic component should interpret the surface string "directly", with no intermediate level of syntax. As one can see from this abstract description, this is not really a different type of compositionality: it is rather a viewpoint as to what the (syntactic) item is that will be compositionally interpreted. After all, many of the practitioners of linguistics that has a level of LF, claim to do a functionally compositional interpretation of this structure. So the opposition is really to the idea that one should be able to manipulate the "underlying" and "hidden" syntactic structure of natural language in such a way that a function can be described that will generate a semantic evaluation in a compositional manner.

This non-LF way of emphasizing the organization of the grammar is aided by the view of the grammar's rules taken in the Direct Compositionality literature . In this literature, grammatical rules are seen as operating on triples of the form <string, syntactic category, meaning>. And with this, it is easy to see any rule as potentially operating on all three components at once – or rather, having simultaneous effects on the phonology (the string), the syntactic structure, and the meaning.

Jacobson (2012) (slightly modifying the description from her Jacobson (2002)) describes four "levels" of compositionality, and is concerned to show that none of the argumentation in the literature gives a reason not to opt for the "relatively strong" version, Direct Compositionality-2. This is the view where the syntactic structures to be the domain a functionally compositional semantics consists of the surface string but with some accommodation for "dissociated elements" along the lines of the 'wrap'-operations of categorial and tree-adjoining grammars. At a minimum, this will require the syntax to "keep track" of the possible locations that allow for strings to be inserted into another string. (E.g., *take out* is allowed to wrap around *the garbage* so as to accommodate *take the garbage out*.) Generally speaking, difficulties for any view in the general realm of string-oriented theories will be about the ability (or not) to identify structural ambiguities ("bracketing ambiguities"), together with a concomitant lack of resources to syntactically distinguish the two "brackettings". But then semantic compositionality will yield only one meaning, and will not have the resources to distinguish the semantic ambiguity.

8.4 Semantic Groundedness

Pelletier (1994b) introduced the notion of semantic groundedness as a type of semantic atomism that was not compositional. The idea was that compositionality had a certain similarity to inductive definitions, and that once we see that there is a way to generalize those sorts of definitions, we would see a similar way to generalize semantic compositionality. It would no longer literally be the sort defined in Definition (1) as functional composiitonality, because it was not restricted to interpretations of a complex unit that were functions of that unit's parts. Examples help explain the conception.

The relevant aspect of an inductive definition of a function $\boldsymbol{f}$ is the way such definitions always defined a "larger case", $\boldsymbol{f}^{n+1}$, in terms of that same function applied to some "smaller case(s)" and the result of that application was then operated upon by some other function. Common mathematical illustrations are the definitions of *factorial* and *exponentiation*:

$$\begin{array}{lll} n! = 1, & \text{if } n = 0 & \mid \quad n^m = 1, \quad \text{if either } n = 1 \text{ or } m = 0 \\ \quad = (n-1)! \times n, & \text{otherwise} & \mid \quad = n^{(m-1)} \times n, \quad \text{otherwise} \end{array}$$

In these cases we see that $n!$ and n^m are defined by employing that same operation to a smaller number and then "doing something" to it (here, multiplying by n). In all such cases, the definition is always guaranteed to "ground out" into the basis case(s), due to the particular constructions always getting "closer and closer" to the basis case..

But there can be a generalization to this sort of definition where the inductive clause is replaced by a clause that does not always operate on something "smaller than" or "a part of" the item currently being defined. Nonetheless, sometimes such a definition will "ground out" in the basis cases. An example of Kaplan (1968/1972)'s (which he used for a very different purpose) can make the point simply:

> Consider propositional logic as formulated with the set of connectives $\{\neg, \wedge, \rightarrow, \leftrightarrow\}$, all of them taken as undefined primitives. Now, let $\boldsymbol{f}$ be an assignment of truth values to atomic sentences.

Then: $\boldsymbol{f}^*(\phi) =_{df}$:

$$\begin{aligned}
\boldsymbol{f}^*(p) &= \boldsymbol{f}(p), \text{ if } p \text{ is atomic} \\
\boldsymbol{f}^*(\neg\psi) &= 1, \text{ if } \boldsymbol{f}^*(p) = 0 \\
\boldsymbol{f}^*(\psi \rightarrow \theta) &= 1, \text{ if either } \boldsymbol{f}^*(p) = 0 \text{ or } \boldsymbol{f}^*(p) = 1 \\
&= 0, \text{ if both } \boldsymbol{f}^*(p) = 1 \text{ and } \boldsymbol{f}^*(p) = 0 \\
\boldsymbol{f}^*(\psi \wedge \theta) &= 1, \text{ if both } \boldsymbol{f}^*(p) = 1 \text{ and } \boldsymbol{f}^*(p) = 1 \\
&= 0, \text{ if either } \boldsymbol{f}^*(p) = 0 \text{ or } \boldsymbol{f}^*(p) = 0 \\
\boldsymbol{f}^*(\psi \leftrightarrow \theta) &= 1, \text{ if } \boldsymbol{f}^*(\psi) = \boldsymbol{f}^*(\psi) \\
&= 0, \text{ if } \boldsymbol{f}^*(\psi) \neq \boldsymbol{f}^*(\psi)
\end{aligned}$$

This definition of $\boldsymbol{f}^*$ is completely compositional. One defines $\boldsymbol{f}^*$ applied to any formula in terms of how it is applied to the (immediate) parts of that formula. But consider replacing the final portion of the definition with:

$$\boldsymbol{f}^*(\psi \leftrightarrow \theta) = \boldsymbol{f}^*((\psi \rightarrow \theta) \wedge (\theta \leftrightarrow \psi))$$

Now we have defined the $\boldsymbol{f}^*$ when applied to a certain category of formulas (the $\leftrightarrow$ formulas) in terms of an entirely different formula that has only a very tenuous relationship, syntactically speaking, to the $\leftrightarrow$ formula. (Recall that all these connectives were taken as primitives.) Still, we all know that this is a correct definition, and that it will work. As can be seen, this new definition is non-compositional. It defines the $\boldsymbol{f}^*$ function when applied to the $\leftrightarrow$ in terms of things that are not a subpart of the double arrow formula.

Any such definition will nonetheless be correct if it ultimately "grounds out" in some basis cases. So, the functions employed in functional compositionality need not be ones that operate on the "parts". They can in theory be *any* functions – so long as they ultimately ground out. This is therefore an acceptance of the Atomistic picture: there are ultimate simples (in the case under consideration) and complex items are defined as functions applied to them. But it denies that the only way to arrive at these atomic elements is by means of the simple "inductive functions" presumed by classical semantic compositionality. These other functions were called "general recursive functions", and one can perhaps see a commonality between this sort of non-compositionality and that described in (Pagin & Westerståhl 2010: pp. 253-4, fn. 8) and (Pagin & Westerståhl 2010a: §3.2), where they complain that Potts (2007)'s account of quotation is not compositional, but is instead recursive.

More generally, the picture of semantics in the "semantic groundedness" view allows a semantic evaluation to bring into play all kinds of facts, all kinds of information; it could bring in context, it could bring in larger stretches of discourse, it could bring in inferences, it could bring in world knowledge to evaluate an expression, where these facts, etc., are not

part of the meanings of the parts of the expression (nor of the method of combination). So long as this evaluation is always grounded – or perhaps more weakly, that it is grounded in the cases where we actually employ it – then this will serve just as well as any compositional analysis. Furthermore, it can account for the "arguments in favour of compositionality" – in fact, it gives essentially the same rationale as compositionality. Another way of making this claim is to say that the arguments in favour of compositionality in §10 are actually arguments in favour of semantic atomism – but semantic compositionality is just one form of semantic atomism, and thus is not the only view that the arguments support.

9 Some Views of Meaning Seemingly Opposed to Semantic Compositionality

There are various established views of what constitutes "meaning" that are commonly seen as contrary to a compositional framework; and in fact, many of the advocates of each of these views are pretty vociferous in their denial of semantic compositionality, although I will make an attempt at showing that at least some of them might be adopted into a compositional theory. The first first group of views might be called Subjectivism (about meaning). One variant on this is Gricean speaker meaning. But another is held by at least the mainstream Cognitive Linguists (and maybe by all of them). To the extent that this latter variant is of the Atomistic variety, it might be given a "recursive semantics" or sometimes perhaps even a compositional semantics. A second group, which can perhaps be seen as some sort of philosophical generalization of Subjectivism (but which is nonetheless importantly different), is called semantic (or meaning) holism. And a third group – which is also sometimes called semantic holism but is better called Fregean contextualism – which in my opinion is *not* opposed to semantic compositionality and in fact can be seen as itself fully compositional despite its initial appearance. I discuss all three of these views in this section, with an eye to determining whether any of them really can accommodate a compositional semantic treatment.

9.1 Subjectivism and Objectivism

Most philosophers of language are Objectivists. They think of meaning in terms of describing "the world", or "the way(s) the world might be". This can be seen by considering such philosophical views as truth-theories of meaning, or possible world theories of meaning, or truth-in-a-model theories of meaning, or information-theoretic theories of meaning. The overarching idea is of the meaning of a piece of language as describing or identifying or putting limits on how the world might be, so as to "fit" with the linguistic item under

consideration. A number of objectivists use the terminology introduced by Frege of *concept* being an "external, objective, but abstract" entity that exists entirely outside the minds of speakers, and use *thought* to designate the same sort of "external, objective, but abstract" entity that corresponds to an entire sentence. A similar division comes into semantics with the distinction between *intention* as a (subjective) mental act or activity that individuals have when they have some goal or purpose, versus an *intension* as an objective semantic value. However, these usages go against the more "normal" usage of the Subjectivists. In Subjectivism, all these terms denote notions that are mind-dependent entities, dependent on the mind of a thinker who is having or possessing them. They are private – only one person has any particular one of these objects.

9.1.1 Gricean Subjectivism

Grice (1957) and other works of his collected in his (1989) (see also Schiffer (1972)) introduced the notion of *speaker meaning* as the basic meaning-starting point, with such notions as timeless or conventional meaning as concepts to be defined in terms of the basic speaker meaning.

Although there have been various iterations in the definition of speaker meaning, the underlying idea is always that a speaker *intends* to convey some message (a belief, or a desire for action, or ...) to an audience by means of that audience's recognizing this intention of the speaker. (We won't follow up on the various refinements of this.) We can see, though, that even sticking with just the simple case where the speaker wishes a hearer to come to have a certain belief, the interpolation of the speaker's intention will, in the general case, be able to divorce itself from what the hearer is able to recognize as the speaker's intention. That is, because it is speaker-meaning which is basic in the Gricean picture, and so the fundamental notion will be the speaker's various intentions, there can be no guarantee of a *function* that will generate a unique meaning from the sub-intentions relevant to the parts of any intended message. (Even assuming there are such sub-intentions – which is not at all obvious in the first place.) In particular, the Grice picture places individual cases of sentence-meaning as the primary basis of meaning. This suggests that there won't be the possibility of the entirety of a speaker's language – all of his/her sentences and their relevant intentions – being susceptible to a compositional analysis. It is only when these basic elements actually form a particularly coherent class that a compositional semantics for the language can be given. We look at such a case below in §9.2.

9.1.2 Subjectivism and Cognitive Linguistics

In mainstream Cognitive Linguistics – and also in related areas of psychology such as cognitive psychology and psycholinguistics – the notion of linguistic meaning is identified

with some part of the conceptual structure of a speaker. A word, for example, is taken to mean a *concept*; a phrase is taken to include the result of a conceptual combination of concepts. (In the Subjectivist's meaning of 'concept'.) To the extent that the meaning of entire sentences is discussed, it appears that it means some sort of "judgment" that is generated when the combination of the right sort of concepts occurs.

One can think of two ways that such a subjectivist view of meaning might be thought to be realized. One is most easily conceived as a version of "subjective atomism", where all the concepts that are lexical meanings are each thought to be basic, atomic simples. A slight variant on this would be that only some lexical concepts are simple, and there are other, more complex lexical concepts that are in some way a combination of the simple concepts. The former picture is usually associated with Jerry Fodor (e.g., in Fodor (1998, 2000, 2001), Fodor & Pylyshyn (2015)), where as a part of the theory it is held that *no* concepts are definable in terms of others... and hence *no* natural language lexical terms are definable. This leads pretty directly to semantic compositionality, as Fodor is often at pains to point out. The latter picture is held by those theorists who think there to be a (finite) set of semantic primitives, perhaps along the lines of Osgood et al. (1957) or perhaps works in the Natural Semantics Metalanguage (NSM) school (e.g., Wierzbicka (1996), Goddard & Wierzbicka (1994), Goddard (2002)) or Shankian universals (Schank (1972)). Here the issue of compositionality seems to reside in the "definiteness" of the manners of possible combination. Given the NSM view that these primitives are common to all languages (and hence maybe are innate to all humans?) perhaps this view, like the Fodor view, would be of a semantically compositional mental (as well as public) language.

A different way to realize the subjectivist viewpoint is taken (almost?) universally in the many variants of Cognitive Linguistics that are on the market. Not only does it deny that lexical concepts are atomic, but it embraces a very wide-ranging panoply of items into these concepts. Vyv Evans (2009, p. 61), cites with approval Langacker (1987):

> What Langacker appears to have in mind is that the semantic material – informally the meaning – associated with a lexical form, i.e., a word, relates directly to the contents of conceptual structure. In principle, this conceptual structure relates to a diverse and sophisticated body of non-linguistic knowledge .... The meaning of *uncle*, on this view, is potentially a function of the vast body of encyclopaedic knowledge we have of what it means to be someone's uncle. In addition to the specific relationship holding between the child of uncle's sibling, this also includes detailed knowledge relating to marital relations, familial relations, the social status of uncles, the types of behaviours associated with uncles, as well as individual knowledge any given individual may have with respect to uncles they have known. Yet while this knowledge is encyclopaedic, it is for Langacker part of semantic

> structure, i.e., directly encoded by a lexical form. Langacker's argument is that there is no principled way of separating putative linguistic from non-linguistic semantic representation.

While there is a sense of subjectivism that amounts to being speaker-dependent, in the Langacker conception of meaning as just described, it is "private" to the speaker – unlike the case for Objectivists – and thus there is also a need for some notion of "similar mental structures" so that communication actually takes place between speaker and audience. And especially challenging for this conception of meaning is the emphasis placed an the individual variation in meaning that any term of public language is alleged to have, in the minds of different speakers/hearers. It has seemed to many that there can be no empirical evidence for actual similarity of such concepts, other than the question-begging belief that people actually do *seem* to understand one another. But the theory itself belies this claim: the theory *emphasizes* the manifold differences between each language speaker and the next.

It should be obvious that the first conception of Subjectivism – the one associated with Fodor – allows for a compositional semantic theory that is common to all speakers of the language. (There are, naturally, difficulties in explaining how the atomic concepts should come into existence uniformly across all speakers of the same language. The standard understanding of Fodor is that he holds these to all be innate. Not only innate within one language, but across all languages: an innate endowment of humans. Not everyone is comfortable with this view.) But the second conception is much more difficult to make into such a theory. Even in the most favourable of circumstances, where a speaker just *happens* to get experiences that make all the different concepts be able to combine compositionally so as to uniformly and systematically yield a result for meanings of complex combinations of the speaker's concepts that can be described in a compositional way, it is pretty clear from the description of the process outlined in the quotation from Langacker that it will not at all be the same function that is involved with any other speaker of the language; so it fails the communicative test. And indeed, (almost?) all the Cognitive Linguists in fact deny that meaning is compositionally organized. Not, it should be added, that any of them have given much thought at all as to how the meanings of the range of complex meanings might be generated, other, perhaps, than as a hodgepodge of "metaphorical extensions". The level of complexity seems to be stalled at individual words – as in the quote from Langacker – and in "conceptual combination" of two concepts. And even in this case they seem not to discuss cases such as simple noun-concept plus verb-concept combinations that would yield a complete thought – a "judgment", such as expressed by *Dogs bark* – as opposed to just a more complicated noun-like concept such as MOUNTAIN LION or BARKING DOG. (For a further description of the areas where the current theories of this sort have been incomplete, see (Pelletier forthcoming); for some actual work on these areas, see especially various of

the works of James Hampton and co-authors, such as his [1979, 1982, 1987, 1988, 1997, 2012, forthcoming].)

9.1.3 Subjectivism and Conceptual Role Semantics

Conceptual role semantics shares certain features with the Langacker view of meaning as "all the idiosyncratic information that a speaker might have concerning a concept". The difference is in the way meanings of individual words are seen as connected to meanings of other words – or to put it in terms of mental concepts, the semantic content of a concept contains as parts all the *liaisons* it has with other concepts. (Liaisons in this picture involve inferential connections, connections by means of opposition, and more generally, any sort of connection between concepts that makes one concept involve – in any way – another concept.) In the Langacker-style view, even though any given concept has much encyclopedic information, each concept is nonetheless seen as an "independent entity" in its own right. . . just a very much more complex and complicated entity than the Fodor-style concept. Conceptual role semantics is an extension of the Langacker view, where instead of viewing each concept as being "self-sufficient" (despite its containing information involving all sorts of other concepts), there is explicit connection made to other concepts in such a way that the connection to these other concepts becomes a part of the meaning of the concept under consideration. Thus, while on Langacker's view UNCLE contains facts about the social status of uncles in the speaker's society (and the other features mentioned), nonetheless this concept is somehow a semantic unit on its own, and the meaning of 'uncle'.

But on conceptual role semantics this last claim is denied. Instead, it is held that the content of UNCLE *includes* the liaisons or connections with the concept of SOCIAL STATUS (and the other concepts mentioned). But then, since what UNCLE includes are the *liaisons* with other concepts and not merely the features of those concepts, it follows that part of the semantic content of UNCLE will also be all the further liaisons that these other concepts themselves have with still other concepts. As a result, the semantic content of any one concept – its *meaning* – will contain information about all the liaisons that all the concepts involved in the "liaison transitive closure" of that concept. Most people who believe in conceptual role semantics agree that this sort of "transitive closure" amounts to *all* the concepts expressed by all words in the lexicon. And they usually call this picture by the name "semantic holism". Or in other words, meaning now becomes a feature of the entire set or body of concepts, and insofar as one can attribute meaning to any one concept, it is most accurately said that this concept's meaning is its place in the entire network of liaisons among the concepts.

Many theorists – and not just Objectivists – find meaning holism wanting ((Fodor & Lepore 1991) is perhaps the best-known opponent). First, as with the Langacker-oriented Subjectivists, because it would be very likely that different people will have had different

language-learning experiences, and therefore will have different liaisons among their lexical items, it follows that they do not speak with the same meanings associated with their words as one another. So, it follows that they do not mean the same thing as one another even when they utter the same word, phrase, or sentence. But then, when they (think they) agree with one another, or disagree with one another, they are in fact not doing so; at the very least, any evidence you have that you disagree with your conversational partner is, according to the theory, better evidence that you are talking about different things, and not disagreeing at all. And should you (appear to) *agree* with someone whom you don't like, that also becomes evidence that you are talking about different things.

But a further feature of this theory's presuppositions, as against the Langacker-style theory, is that not only is communication with others an impossibility, but one cannot communicate with oneself from one year to the next, one day to the next, one hour to the next, ... For, in the intervening time period the person would have new experiences (except in very special cases where the person is comatose) and these will impact the class of liaisons that impinge on the meanings of his lexicon. As opposed to the Langacker view where new information will be encapsulated into some smallish set of concepts and not afflict areas of the lexicon that are not directly involved, in the meaning holism view, *any* change *anywhere* in a person's mental life will change the meanings of *all* concepts. Although one *thinks* one remembers that the bark of eucalyptus trees peels off yearly, there is no justification for this, at least according to the theory. For, over the year when you last had that thought, many of your liaisons have been altered. For example, perhaps you learned that eucalypts exude a substance that is toxic to other plants, and this makes them easier for koalas to climb. So now your your last-year-thoughts have no necessary connection with this-year-thoughts; after all, this has changed your eucalyptus concept contains much new (encyclopedic) information about toxins and koalas (which concepts have also changed) so that now you might think that you were thinking of some other species of tree when you thought that their bark peeled.[10]

Indeed, this sort of meaning-holist can't change his mind about *anything*! For, the very act of getting new information makes it be a different thing, and hence it is *not* a case of changing one's mind about concept X. The old concept X is no longer there to have a new opinion about. Not only is it impossible to disagree with another person, as I remarked above, but it is impossible to disagree with the past. We think that we have learned that Anaxagoras was wrong when he said that the unevenness of the moon's surface is due to the mixture of earthy matter with cold. But this semantic holism is committed to claiming that we are in fact not at all disagreeing with him. We can't disagree. We can't agree either. We've just changed the subject.

Note that this theory too – like the Langacker-style theory described above in §9.1.2 – will want to avail itself of some notion where the "conceptual frameworks" of different

people (or of one's own self at different times) are "similar enough" to allow for communication. In the case of the Langacker-style theory, I said the notion of "similar enough" was merely a question-begging re-description of the phenomenon. The same goes here.

As argued above in §7.1, even though not all Atomistic theories are necessarily functionally compositional, *only* Atomistic theories stand a chance of being functionally compositional. Conceptual role semantics is not a form of Atomism, and hence cannot be compositional. It is instead a form of holism; so there are, in this theory, no basic or self-subsisting entities at all. Rather, one must fit the whole of language into some larger whole, so as to understand *any* part of language. For instance, it must be fit into a theory of society, or culture, or (heaven help us) a Form of Life.

9.2 Fregean Contextualism

Frege (1884) says the following:

- Never ask for the meaning of a word in isolation, but only in the context of a sentence.
- It is enough if the sentence as a whole has a sense; by this also its parts get content.
- We ought always to keep before our eyes a complete sentence. Only therein do the words really have a meaning.

Scholars dispute the meaning of these rather unexplained claims (see, for example, Janssen (2001, 2012), Pelletier (2001) and the many references therein), but one natural way is to see them as advocating an *anti*-atomistic view of meaning, thinking that sentences are to be given a priority in theorizing about semantics, because of the priority of judgments or ("complete") thoughts, and so on. All other meanings in some way supervene on these primary meanings. And many researchers have taken these remarks in that way.

But on the other hand, they could be taken as instead making a claim about the "total extent" of the meaning of any lexical item – that is, as saying that the totality of all the sentences that use a word exhausts all of what that word can mean – then we can see it as a constraint on what synonymy is (Hodges 2001). Some accounts of at least some words do take this view already. Syncategorematic words – such as quantifiers are thought to be, on some accounts – do not have any independent meaning other than "what their effect is in a sentence". To discover the meaning of *each*, for example, as it occurs in sentences like

(24) a. Each Canadian citizen over 21 years of age is permitted to vote.
b. Bill ate each of the biscuits slowly.
c. The company guarantees the quality of each item that anyone purchases.

we compare their meanings with the meanings of sentences where *each* is replaced by other quantifiers. This might be said to give the full extent of the meaning of the word *each*. (There is of course much more to it than this: among the sentences to be compared are ones with negations, so somehow the meaning effect of *each* on the sentence will be conditioned by the meaning-effect of negation. As a practical way to determine the meaning of words, even just some small group of syncategorematic words, this use of Fregean Contextualism is decidedly difficult. But that is not Hodges' point here, which is to consider the effect of *constraining* the meanings of any word to be the ones which are manifested in some sentence or other.)

In this setting, Hodges proposes his "Husserlian Property"

Definition 10. (Husserlian Property) *Synonymous terms belong to the same semantic category. That is, synonymous terms are intersubstitutable in any complex term containing one of them, without loss of meaningfulness.*

The following corollary is almost automatic from this:

Corollary 1. (Meaning Constraint) *If two terms, a and b are* not *synonymous, then there is some meaning difference between some larger term that contains one of a and b (say a), and that same larger term where a is replaced by b. This can happen in two different ways: either exactly one of the two larger terms is meaningful or else they are both meaningful but are not synonymous.*

Given the Husserlian Property, this Meaning Constraint, and either functional or substitutional compositionality, it is again almost automatic to get

Corollary 2. (Fregean Contextual Meaning) *The meaning of a term is the contribution it makes to the meanings of complex terms (including sentences) of which it is a part.*

This way of approaching Fregean Contextual Meaning implies that for every non-synonymous pair of terms, there are some sentences that differ only in containing one vs. the other of this pair, and these sentences are not synonymous. It amounts to saying that if two terms make the *identical* contribution to *every* complex term, then their meanings cannot be distinguished. Seen this way, Fregean Contextual Meaning is not in conflict with semantic compositionality. Indeed, it is more or less the same thing.[11] (See Pagin (1997, 2006) for further philosophical thoughts on this topic. The formal material can be found in Hodges (2001).)

10 The General Arguments for Semantic Compositionality

A common claim in the formal semantic literature is that a compositional semantics is "neat and tidy", that it gives direction to the construction of a full grammatical theory, and that it makes it clear how "language relates to reality." More abstractly, we can say that such claims emphasize certain aesthetic features of compositional semantic theories, such as its clarity and transparency. But while many formal semanticists do hold this to be an important advantage of such theories, certainly if they acknowledge that there are any relevant empirical data at all, then they would say that such data should overcome this feeling of aesthetic pleasantness. Certainly aesthetics doesn't seem to be a *direct* argument for semantic compositionality. (Not every theorist thinks there is any relevant empirical data. Some hold that semantic compositionality is a *vacuous* claim with no empirical content: that *any* language at all can be given *some* compositional semantics. We discuss this claim below in §11.2.

Besides this "argument from aesthetics", there are three (or four, depending on how one individuates arguments) "arguments", related to one another, that are traditionally given to support semantic compositionality. Frege (1923) apparently gives a version of at least one of the arguments:

> It is astonishing what language can do. With a few syllables it can express an incalculable number of thoughts, so that even a thought grasped by a human being for the very first time can be put into a form of words which will be understood by someone to whom the thought is entirely new. This would be impossible, were we not able to distinguish parts in the thought corresponding to the parts of a sentence, so that the structure of the sentence serves as an image of the structure of the thought.

Similar sentiments can be found in other authors of the same period.[12] Generalizing from these considerations of Frege's, we can discern three variant versions, focussing on slightly different features of "language". (It is admitted by everyone concerned that these are all "arguments to the best explanation" and do not *logically* entail the conclusion of semantic compositionality.)

Argument from understanding We can *understand* an infinite number[13] of novel sentences, so long as they employ words we already understand. We understand sentences and other combinations that we have never encountered. So, language must be 'compositional': it must start with a finite stock of words/morphemes with their meanings, and put these together in a finite number of different ways, but

using an unlimited recursive method to arrive at the infinite number of understood sentences.

Argument from productivity/creativity We can *create* new sentences that we have never heard or used before, and we know that they are appropriate to the situation in which we use them. This can only happen if language is 'compositionally' organized, so that we learn some finite base of words and rules, but know how to combine them recursively so as to produce totally new descriptions with the intended meanings.

Argument from learnability We are finite creatures who are exposed to a finite amount of information concerning our language. Nonetheless we can *learn* a system that is capable of infinite expression of meanings. The only way this can happen is if what we learn has a finite basis of terms and rules, but the rules themselves allow for arbitrarily complex 'composition'.

As I see it, there are three threads interwoven in the arguments:

i. that language is something special (infinite, or novel, or creative, or whatever) in its ability to express meanings,

ii. that people manage to use/learn/understand language despite their being "finite",

iii. that one (the only known?) way to do this is if language exhibits a semantically compositional framework.

Many authors have found these arguments convincing, not least because there seems to be no alternative account of the various abilities. But it might be noted that in §8.4 I rehearsed the considerations from Pelletier (1994b), where it was alleged that these arguments really only justify some form of semantic atomism, and not necessarily the particular form that is taken by semantic compositionality. In particular, Pelletier argued that Semantic Groundedness would be enough, and that this version of atomism did not always follow the literal statement of Definition (1).

Other writers have voiced other worries about the cogency of the argumentation. Szabó (2012: pp. 74–77) worries about the presumed notion of understanding that is inherent in the arguments: that understanding is a matter of grasping the meaning of that which is understood. He puts the case thus:

> This is certainly not right in general: to understand a problem, an idea, or a proof is not the same as grasping the meaning of the problem, the meaning of the idea, or the meaning of the proof.... [O]ne might try to stipulate that ...to understand a *linguistic expression* is nothing more or

> less than grasping what it means. But such a stipulation is by no means innocent. There are many conceptions of meaning for which it is a rather bad fit. It is certainly false… that understanding a linguistic expression is a matter of grasping its extension.… [Could it] require that we grasp a set of possible worlds? …Does it require the ability to tell of any particular possible workld whether it is a member of the set? …Should [a holder of the direct referential view of proper names] accept that understanding a proper name is a matter of grasping its bearer?

Szabó also has worries about the argument from productivity, saying it assumes "that we *already* understand expressions we have never heard before. What is the evidence for this? The fact that when we hear them we understand them shows nothing more than the information necessary to determine what they mean is available to us immediately *after* they have been uttered."

These concerns are only a tip of the iceberg, since all theories in the end will need somehow to find a link between the mental life of speakers/hearers and the meanings of linguistic expressions.[14] But as the cognitive psychologists and cognitive linguists are often at pains to point out, this seems to involve much more than the notion of "grasping". And all the available (partial) explanations seem also to violate semantic compositionality.

11 (Some) Formal Properties of Semantic Compositionality

11.1 Semantics as a Homomorphism of Syntax

Almost every formal description of semantic compositionality, and a number of the others, remark that semantic compositionality views the semantic component of a grammar as a homomorphism of the syntax. This follows the line of thought initiated in Montague (1970: §§1, 3), where both the syntax of a language and the semantics of a language are described as algebras, and the semantic algebra as a homomorphism of the syntax. (A very nice, gentle introduction to this is given in Janssen (1997: pp. 448–450).) A homomorphism from an algebra $\mathscr{A}$ to another algebra $\mathscr{B}$ is a mapping which respects the structure of $\mathscr{A}$ in that, whenever an element a of $\mathscr{A}$ is generated in $\mathscr{A}$ by applying an operator f to other elements $a_1, \cdots, a_n$ of $\mathscr{A}$, then the image in $\mathscr{B}$ correlated to a is obtained by the operator of $\mathscr{B}$ that is correlated with f, as applied to the elements of $\mathscr{B}$ that are correlated with $a_1, \cdots, a_n$. That is to say:

Definition 11. (Homorphism) *If h is a homomorphism from algebra $\mathscr{A}$ to algebra $\mathscr{B}$, then whenever an item a in $\mathscr{A}$ is generated by operator f applied to $a_1, \cdots, a_n$, then there is a unique item b (namely $h(a)$) in $\mathscr{B}$ which is generated by a unique operator $h(f)$ (the correlate of f) applied to the items in $\mathscr{B}$ that are correlated with $a_1, \cdots, a_n$, that is, applied to $h(a_1), \cdots, h(a_n)$.*

Besides differences in the names of their elements, structural differences between $\mathscr{A}$ and $\mathscr{B}$ will therefore be due to two facts only: that two different elements of $\mathscr{A}$ can be correlated with the same member of $\mathscr{B}$, and that two different operators of $\mathscr{A}$ can be correlated with the same operator in $\mathscr{B}$.

So, in the case of semantics of a language, the syntactic language corresponds to $\mathscr{A}$ above and its elements (called a above) are the lexical items and other syntactic terms of the syntax. The operations are the syntactic rules of the language, and correspond to the f above. The mapping to a semantic language is homomorphic when, to each of the as in the syntax, there is some unique element or image in the semantics that corresponds to it; and for each syntactic operation f that employs $a_1, \cdots, a_n$ and creates some syntactic element a_{n+1}, there is some unique semantic operation in $\mathscr{B}$ that relates the images for each of $a_1, \cdots, a_n$ to the unique image of a_{n+1}. As remarked, it is possible that even though $a_1, \cdots, a_n$ are all distinct, not all of the semantic images of them differ from one another. And it is possible that different syntactic operations in $\mathscr{A}$ can correspond to the same semantic operation in $\mathscr{B}$. In a way, then, the semantic language "loses information" about distinct items of the syntax, but then perhaps "gains information" by asserting the semantic identity of distinct items of the syntax.

This balance of "loss" vs. "gain" of information forms the basis for certain objections to semantic compositionality that are presented above in §6.4.

11.2 Vacuity

There have long been worries about the possible "vacuity" of semantic compositionality – that is, worries as to whether maybe it is a meaningless condition because *any* language could be given a compositional semantics. If that were the case, then compositionality would be only a matter of taste or methodological preference, and not an empirical claim. (Recall the discussion in §5.1 to the effect that the notion of a function is so underconstrained that all sorts of things could be gerrymandered into the semantic representation of any construction.) A proof in Janssen (1986) showed that every set of strings that can be generated by a Turing Machine – thus, every recursively enumerable set – can be given a compositional semantics. Although some have found this proof to be important, it should be noted that the method employed to show this is allowed to freely invent any grammar whatsoever and to freely assign any semantics whatever to this invented grammar. So, the

claim really amounts to an assertion that every recursively enumerable set of strings has *some* grammar that admits a compositional semantics. The total freedom to construct *any* syntactic analysis with no constraints probably doesn't fulfill what most theorists would want to call "constructing a compositional semantics for an arbitrary grammar". Probably most theorists think of the task as one of constructing a compositional semantics *from a given syntax for the language*, and that cannot always be done.

Besides the Janssen (1986) work just described, other earlier work by van Benthem (1986), Hendriks (1993) also contained results that some writers have interpreted as making the same general sort of claim – that it is always possible to construct a compositional semantic theory. The issue came to a more formal head with the appearance of Zadrozny (1994), where a construction was given that converted any language that had any semantic theory into one that had a compositional semantic theory and from which the original meanings that were generated by the non-compositional theory could be recovered.

Zadrozny's article generated many responses, such as Westerståhl (1998), Kazmi & Pelletier (1998), Dever (1999), mostly finding flaws in the inference from the mathematical construction offered to the conclusion that he has given a compositional semantics for the language in question. Kazmi & Pelletier (1998) claim that Zadrozny's result is not a theory of meaning in the original sense. At most it is a compositional theory of "pointers to meanings", they say, but it is not true that the meaning, in the original intuitive sense, has been used to construct a compositional theory of meaning – it is not true that the original semantics has been encoded as a compositional semantics. Westerståhl (1998: pp.641–642), who was considering also the Janssen and van Benthem works mentioned in the previous paragraph, in addition to Zadrozny, says

> . . . [These results] achieve compositionality by syntactic and semantic manipulations that are of no interest to the semanticist. Janssen's theorem introduces a completely ad hoc syntax. Automatically extending a mapping of lexical items to a homomorphism as van Benthem suggests violates basic intuitions in actual examples. Zadrozny's theorem, the most striking result in this vein, makes the meaning assignment one-one in an unmotivated way, thereby side-stepping the compositionality issue.

So, it seems that the vacuity occurs only if one is allowed not only to re-jig the semantics, but to change various aspects of the syntax also. Surely this isn't what is meant by saying that "every semantics can be recast as a compositional semantics"!

The commentary in Dever (1999) is more far-reaching, questioning the formal development as well as the interpretation that Zadrozny gives to his results. Of interest to us is Dever's claim that "one rough constraint [on a new meaning function] is that [it] respect our independent judgments about the meanings of the parts. . . . Of course,. . . it seems

acceptable that there be some regimentation of semantic values. . . Nevertheless, our starting intuitions must be recognizeable in the final result." One constraint Dever imposes is that the new compositional meaning theory has to assign the same semantic value to terms we initially took as synonymous. Dever also distinguishes portions of the newly constructed meaning function that are constrained by the original meanings (the "manifest meanings") from those which are not, but which instead tell how the syntactic terms act when combined with other terms in the language. These latter are called "Occult Meanings" by Dever and he finds that

> the occult is all-powerful – the occult meanings in Zadrozny's system are doing all the work.. . . At every stage (of a derivation or conversion to a compositional semantics) the manifest meanings of the terms are inert, while the occult meanings fully determine the value at the next stage.. . . We could change every manifest meaning in the system, and. . . we would not alter in the slightest the computational procedure (or its result) through which we determine the value of wholes from their parts. Since only the manifest meanings of terms are constrained by the original meaning function, we [are provided with] only the weakest of connections with our original intuitions on meaning. (Dever 1999: p. 318)

Dever closes this article with a consideration of what would be the effect of requiring constraints on various parts of the language; for instance, thinking that the pre-theoretic view of the meanings of sentences is primary and must be maintained in the newly-constructed compositional theory but that other parts of grammar are epiphenomenal. So a compositional theory would be required to maintain the same meanings for sentences, even though they can employ occult meanings for other parts of language. Dever shows that a compositional theory can always be constructed this way (although it requires some unusual way to hide the occult meanings when it comes to computing the final sentence meanings), and claims that this "does not allow for the use of compositionality as a strong constraint on theory formation." On the other hand, if we require that the lexical items maintain their pre-theoretic meanings, then even if we allow for occult meanings at other syntactic levels, "we will not always be able to obtain compositional meanings functions. . . . But if we. . . allow even one atomic term in the grammar to carry occult meaning, then compositonality loses force in any sentence in which that term appears."

11.3 Compositional Extension

The extension problem is this: given a grammar for which there is a partial semantics (partial in the sense that it does not assign a meaning to every term in the syntax) that is semantically

compositional for the portion that it does assign meanings to, can it always be extended to a compositional semantics for the entire language? The first formal investigations about this question were carried out by Wilfrid Hodges in a series of papers that culminated in his Hodges (2001) – see also Hodges (2012a,b). In the earlier papers he was concerned with the case where the set of sentences were given a compositionally semantic interpretation in terms of their subparts but the remainder of the syntactic elements (including these subparts) were not necessarily given such a semantic interpretation. (Thus, one of the consequences of this work was an answer to the Fregean Contextualism issue discussed above in §9.2.)

The setting under consideration is where some arbitrary portion of the language (described in terms of the syntactic terms of the language) is given a compositional semantics. And we wish to know whether *without a change in the syntax of the language* (including the restriction that we add no new words or expressions to the language) we can give a compositional semantics for the whole language. In this setting, Hodges asks for conditions that will allow for one always to be able to "extend" the compositional portion of the semantics so as to hold for the meanings of all syntactic items. In Hodges (2001) he proves that under two conditions on the initially given semantics (in addition to the partial compositionality), it can always be done. The two conditions are *co-finality* (any term not already compositionally meaningful must be a subterm of some already compositionally meaningful term) and the *Husserl condition* (that all terms that are synonymous must be in the same semantic category – see Definition (10) in §9.2 above). With these conditions, not only will there be a compositional semantics for the entire language, but also it will be of a certain sort, called a *Fregean extension*.

Westerståhl (2004) considers some variations on Hodges' results: what would happen if the co-finality condition were dropped, or the Husserl property were dropped, or both? If both are dropped without replacement, then there need not be any total compositional extension. But if the two are replaced with the assumption that the subterms of any meaningful term are also meaningful, then a compositional extension will exist (Westerståhl's Theorem 12).

If we keep the Husserl property on the initial compositional semantics but not co-finality (also not the closed-under-meaningful-subterms property), then by Westerståhl's Corollary 11 it can be shown that there is a compositional extension (which, however, may no longer be Husserlian). On the other hand, if we keep co-finality (and compositionality) of the initial partial semantics, but do not have the Husserl property, Westerståhl offers us only an open problem: it is unknown what further property will be necessary to guarantee a total compositional extension. (Rather, unknown whether there are any more minimal conditions than the Husserl property.)

12 This Chapter

We started this journey through the compositional landscape with the standard definition of compositionality, functional compositionality, and some quick explanations of the terms that are used in the definition. The explanations of some of these terms seem to bring with them (what we called) "theory-internal problems", where it seemed that such notions as *function* would fail to capture what was wanted by compositionality. We also looked to some of the well-known "empirical problems" in the form of natural language constructions that challenged compositionality. The challenges discussed here are just a small fraction of the challenges in the literature, but they do include what some find to be the most difficult to bring under the compositional framework: context, idioms, and quotation; and we saw the sorts of moves made by advocates of semantic compositionality in response to these challenges.

We also looked at variants of the definition, at certain mathematical correlates of the concept of semantic compositionally, and explored how semantic compositionality fits into a broader spectrum of Atomistic vs. Wholistic/Holistic theories. The "arguments" in favour of semantic compositionality seem equally to be arguments for *any* Atomistic theory of meaning, and therefore one can take these arguments seriously while nevertheless denying the letter of statements of semantic compositionality. On the opposite side, non-Atomistic theories will not only find these arguments difficult to deal with, but standard developments of such theories are subject to a wide range of other criticisms.

It is certainly not clear just where the final conclusion will be to this dispute. But the topics engaged in this chapter should show that the answer is not as obvious as either of the two attitudes mentioned in the Summary section at the outset would have you believe.

13 Further Readings

- Partee (1984), reprinted in B.H. Partee (2004) *Compositionality in Formal Semantics. Selected papers by Barbara H. Partee*, (Blackwell: Malden, MA), pp. 153–181. See also her "Genitives: A Case Study" in the same volume, pp. 182–189, which was not in the original published version but was nonetheless written prior to that version of the paper, and at one time viewed as a part of the original paper. Details in the footnote on p. 182 of the latter article.

- T. Janssen (1997) "Compositionality" in J. van Benthem and A. ter Meulen *Handbook of Logic and Language*, (Elsevier: Amsterdam), pp. 417–473. A magisterial overview of the topic as it was in the mid-1990s. It includes the Partee "Genitives: A Case Study" on pp. 464–470. . . the first time it was published.

- P. Pagin and D. Westerståhl, eds. (2001) *Journal of Logic, Language, and Information: Special Issue: Compositionality* vol. 10/1. This special issue contains six articles plus an introduction, and all of them have attained a special place in the area of compositionality. The entries tend to be on the detailed, formal side.
- J. Fodor and E. Lepore (2002) *The Compositionality Papers* (Oxford: Oxford UP). A collection of papers on the topic of compositionality taken from the viewpoint of one of the most vociferous proponents of compositionality in people's (linguistic-oriented) mental life.
- J. Dever (2006) "Compositionality" in E. Lepore and B. Smith *The Oxford Handbook of Philosophy of Language* (Oxford: Oxford UP) pp. 633–666. Defines the notion, paying attention to other definitions. Describes various attitudes towards the principle and offers compositional solutions to some difficult phenomena. Ends by describing *why* anyone would want a compositional semantics. An especially long version of this entry is available on his website. (See link in the next section below.)
- Pagin and D. Westerståhl (2010) "Compositionality I: Definitions and Variants", and "Compositionality II: Arguments and Problems" in *Philosophy Compass* vol. 5, pp. 250–264 and 265–282. This is an excellent discussion of the various formal properties of definitions of compositionality (in Part I), and of various linguistic phenomena that can be handled compositionality even though they seem recalcitrant.
- M. Werning, W. Hinzen, and E. Machery, eds. (2012) *The Oxford Handbook of Compositionality* (Oxford: Oxford UP), xviii + 743pp. This is now the standard reference for all aspects of compositionality and its interaction with many different phenomena.
- J. Hampton and Y. Winter, eds. (forthcoming) *Compositionality and Concepts in Linguistics and Psychology* (Berlin: Springer). A volume that engages the differing views of linguists and psychologists on the issue of how compositionality and concepts might work together, or not.

14 Links to Digital Resources

- "Compositionality" in the *Internet Encyclopedia of Philosophy*: http://www.iep.utm.edu/composit
- "Compositionality" in the *Stanford Encyclopedia of Philosophy*: http://plato.stanford.edu/compositionality

(These are both very nice entries on Semantic Compositionality; the entry in the Stanford Encyclopedia is aimed at a somewhat more advanced audience than the one in the Internet Encyclopedia.)

- A hierarchically-organized and annotated bibliography of more than 150 works on the topic of compositionality is available through the Oxford Bibliographies Online resource: http://www.oxfordbibliographies.com/browse?module_0=obo-9780199772810
- The very long version of Dever (2006), mentioned above in the "General Overview Works" section, is called "The Extended Mix" and is available at https://sites.google.com/site/joshdeverphilosophy/

Endnotes

[1]Some – indeed, many – of the thoughts espoused in this chapter were originally expressed in a series of articles by me: Pelletier (1994b, 2000, 2012, forthcoming). And thanks go to all those who engaged (or corrected) me on many of these issues, including a referee for the present paper.

[2]The Oxford English Dictionary cites Smuts (1926) as the originator of 'holism'. In this work 'holism' is defined as "the theory which makes the existence of 'wholes' a fundamental feature of the world. It regards natural objects, both animate and inanimate, as wholes and not merely as assemblages of elements or parts." Smuts claims that evolutionary forces act on species, rather than individuals, with the consequence that "evolution aims at more perfect wholes". The OED also cites Boodin (1939) as the originator of the word 'wholism': "Two conceptions... namely, creative synthesis (or emergence), and wholism (or gestaltism)...".

[3]A general discussion of "context" in formal (and informal) semantics is in (Predelli 2005), which is concerned to defend "the traditional view" (which presumably includes semantic compositionality) from attacks by "contextualists". (Predelli is to a large extent defending Kaplan's view).

[4]"Intersective" adjectives, as they are usually called. Kamp & Partee (1995: p137) cite *carnivorous* as such an adjective. In many writings (including my text here) there are no such adjectives, and they are assimilated to the case of surjective adjectives, described in the next sentences. Besides surjective adjectives, there are additionally 'privative' adjectives like *fake, toy, former, imitation,...* that are interpreted as picking out some (vaguely-specified) group of things that are *not* in the set of things picked out as the meaning of the modified noun.

[5]In (21-a), if 'Quotation' is a syntactic part of ' 'Quotation' ', for example.

[6]Pagin & Westerståhl (2010) cite Quine (1940), Tarski (1936) as examples of name-theories; Tarski (1944), Quine (1960), Geach (1957), Werning (2005) as examples of description-theories; and Davidson (1979), Predelli (2008) as examples of the demonstrative-theory. They also cite Frege (1892), Searle (1969), Washington (1992), Saka (1998), Recanati (2001, 2008) as proponents of use-theories, which we will discuss in the next paragraph. They could also have cited Bigelow (1975a,b) as a type of use-theorist.

[7]Trying here not to pre-judge the philosophical issue of whether meanings are *properties* of syntactic items or are *independent* things in their own right that are related in some way to the syntactic items.

[8]It is not so clear just what is meant by the "purely interpretive" constraint. One thinks that the authors have in mind that the structure of the syntactic item being evaluated cannot be changed. But it isn't so clear that any of the formal semanticists ever wanted to do such a thing, since the functional compositionality definition evaluates only meanings. On the other hand, Zimmermann (2012) discusses "solutions" to compositionality problems that amount to changing the syntactic structure of the item being evaluated. And Jacobson (2002, 2012) inveighs against theories that alter the syntactic form in the construction of a "LF" that then becomes the input to a semantic component. Perhaps this is what Larson & Segal are trying to rule out, although to my mind this is instead a matter of changing the underlying syntax of the language, so that we are *not* doing anything with compositionality here but instead have decided to alter the language – of course, the reason for wanting to change the language is the difficulty in maintaining compositionality. This makes it seem that semantic compositionality for its own sake is driving the rest of the grammar, rather than being a consequence of the independently-motivated syntax and other parts of the grammar. In the terminology of §10 below, this is a consequence of "the aesthetic argument for compositionality",

[9]Of course, sometimes this level of interpretation can be eliminated – for instance, if there is a *function* that always takes the relevant syntactic structure to the level of interpretation and another *function* whose domain is just the items in the interpretive level and whose co-domain is the final semantic value. This is eliminable by function composition, and the intermediate level is thereby nothing more than a convenience.

[10]As I am presenting it, then, the Langacker-style theory is a variety of Atomism about meaning. The concepts might include very many complex features over and above the picture advocated by Fodor, but nonetheless they ultimately "ground out" in a definite set of features. In contrast, the current viewpoint being discussed in the text is "holistic" in that it never "grounds out" except when *all* the concepts involved in the entire conceptual framework are involved. It is not so clear to me that what I've called the Langacker-theory is a stable picture – isn't one tempted to claim that the inclusion of all the encyclopedic knowledge in each concept requires further and further elaboration, until we end up with the non-Atomistic semantic holism? But I won't follow up on this idea here.

[11]Note the difference between Fregean Contextualism, where the meanings of the set of sentences is given, and the Gricean view, where sentences are seen to be "the primary level" (which we discussed above in §9.1.1. With Grice, these sentences are themselves defined in terms of speaker-intentions; in Fregean Contextualism we are just presented with all the sentences.

[12]Russell (1918: Lecture II, pp. 514–517)

> I come now to another point, which applies primarily to propositions and thence derivatively to facts. You can understand a proposition when you understand the words of which it is composed even though you never heard the proposition before. That seems a very humble property, but it is a property which marks it as complex and distinguishes it from words whose meaning is simple. When you know the vocabulary, grammar, and syntax of a language, you can understand a proposition in that language even though you never saw it before.
>
> . . . I should say that the word 'red' is a simple symbol and the phrase 'This is red' a complex symbol. The word 'red' can only be understood through acquaintance with the object, whereas the phrase 'Roses are red' can be understood if you know what 'red' is and what 'roses' are, without ever having heard the phrase before. That is a clear mark of what is complex.

Wittgenstein (1921) has:

> 4.026 The meanings of simple signs (words) must be explained to us if we are to understand them. With propositions, however, we make ourselves understood.
>
> 4.027 It belongs to the essence of a proposition that it should be able to communicate a *new* sense to us.
>
> 4.03 A proposition must use old expressions to communicate a new sense.

Schlick (1918) says

> A cognitive judgment is a *new* combination made up exclusively of *old* concepts. The latter occur in innumerable other judgments, some of which (e.g., the definitions of these concepts) must have been known to us already. . .
>
> Only the primitive concepts and judgments. . . depend on conventions and have to be learned as arbitrary signs. . . . Language, for its part, operates in a fashion similar to the cognitive process.

None of these authors cite any earlier thinkers for this general argument/observation.

[13]Although the usual statement of these claims mention "an infinite number" (and Frege uses "an incalculable number"), it has been pointed out that even with a finite, but large number of sentences there is a requirement of something like compositionality or recursiveness. After all, these facts hold of finite creatures

– us – and with just a modestly large number of sentences in the language we would not have had the time required to learn, or experience, or bring to mind, all of the sentences. See Grandy (1990).

[14]Among other works, one can consult Robbins (2002) for a listing of many different ways to interpret the different arguments and the various ways to interpret mental items. His conclusion is that none of the arguments can be seen as establishing their conclusions, mostly on the grounds of equivocating amongst the various interpretations.

References

Bach, Emmon. 1976. An extension of classical transformational grammar. In *Problems of linguistic metatheory: Proceedings of the 1976 conference*, Michigan State University.

Barker, Chris & Pauline Jacobson (eds.). 2007. *Direct compositionality*. Oxford: Oxford UP.

van Benthem, Johan. 1986. *Essays on logical semantics* chap. "The Logic of Semantics", 198–214. Dordrecht: Reidel.

Bigelow, John. 1975a. Quotation and compositionality, 1. *Linguistische Berichte* 38. 1–21.

Bigelow, John. 1975b. Quotation and compositionality, 2. *Linguistische Berichte* 39. 1–21.

Boodin, John. 1939. *The social mind: Foundations of social philosophy*. New York: Macmillan Co.

Cappelen, Herman & Ernie Lepore. 2012. Quotation. In Edward N. Zalta (ed.), *The Stanford encyclopedia of philosophy*, http://plato.stanford.edu/archives/win2006/entries/quotation/.

Cooper, Robin. 1983. *Quantification and syntactic theory*. Dordrecht: Reidel.

Davidson, Donald. 1979. Quotation. *Theory and Decision* 11. 27–40.

Dever, Josh. 1999. Compositionality as methodology. *Philosophical Review* 22. 311–326.

Dever, Josh. 2006. Compositionality. In E. Lepore & B. Smith (eds.), *The Oxford handbook of philosophy of language*, 633–666. Oxford: Oxford UP.

Evans, Vyvyan. 2009. *How words mean: Lexical concepts, cognitive models, and meaning construction*. Oxford: Oxford UP.

Fodor, Jerry. 1998. *Concepts: Where cognitive science went wrong*. Cambridge, MA: MIT Press.

Fodor, Jerry. 2000. *The mind doesn't work that way: The scope and limits of computational psychology*. Cambridge, MA: MIT Press.

Fodor, Jerry. 2001. Language, thought, and compositionality. *Mind and Language* 16. 1–15.
Fodor, Jerry & Ernest Lepore. 1991. Why meaning (probably) isn't conceptual role. *Mind and Language* 6. 328–343.
Fodor, Jerry & Zenon Pylyshyn. 2015. *Concepts without meaning*. Cambridge, MA: MIT Press.
Frege, Gottlob. 1884. *The foundations of arithmetic*. Oxford: Blackwell. Translation by John Austin, 1950.
Frege, Gottlob. 1892. Über Sinn und Bedeutung. *Zeitschrift für Philosophie und philosophische Kritik* 100. 25–60. Translated as "Sense and Reference" (or "Sense and Meaning", by different translators).
Frege, Gottlob. 1923. Logische Untersuchungen. Dritter Tiel: Gedankengefüge. *Beträge zur Philosophie des deutschen Idealismus* III. 38–51. Translated by R. Stoothof as "Compound Thoughts" and published in *Mind*, 72: 1–17 (1963).
Geach, Peter. 1957. *Mental acts*. London: Rutledge & Kegan-Paul.
Gillon, Brendan. 2007. Pāṇini's *Aṣṭādhyāyī* and linguistic theory. *Journal of Indian Philosophy* 35. 445–468.
Goddard, Cliff. 2002. The search for the shared semantic core of all languages, In Goddard & Wierzbicka (2002), volume 1, pp. 5–40.
Goddard, Cliff & Anna Wierzbicka. 1994. *Semantic and lexical universals*. Philadelphia: Benjamins.
Goddard, Cliff & Anna Wierzbicka. 2002. *Meaning and universal grammar: Theory and empirical findings, vols 1 & 2*. Philadelphia: John Benjamins.
Grandy, Richard. 1990. Understanding and the principle of compositionality. *Philosophical Perspectives* 4. 557–572.
Grice, Paul. 1957. Meaning. *Philosophical Review* 66. 377–388.
Grice, Paul. 1989. *Studies in the way of words*. Cambridge, MA: Harvard UP. Contains Grice's articles on language, and his later thoughts.
Hampton, James. 1979. Polymorphic concepts in semantic memory. *Journal of Verbal Learning and Verbal Behavior* 18. 441–461.
Hampton, James. 1982. A demonstration of intransitivity in natural categories. *Cognition* 12. 151–164.
Hampton, James. 1987. Inheritance of attributes in natural concept conjunctions. *Memory and Cognition* 15. 55–71.
Hampton, James. 1988. Overextension of conjunctive concepts: Evidence for a unitary model of concept typicality and class inclusion. *Journal of Experimental Psychology* 14. 12–32.

Hampton, James. 1997. Conceptual combination: Conjunction and negation of natural concepts. *Memory and Cognition* 25. 888–909.

Hampton, James. forthcoming. Compositionality and concepts. In J. Hampton & Y. Winter (eds.), *Compositionality and concepts in linguistics and psychology*, London: Springer.

Hampton, James & Martin Jönsson. 2012. Typicality and compositionally: The logic of combining vague concepts. In M. Werning, W. Hintzen & E. Machery (eds.), *The Oxford handbook of compositionality*, 385–402. Oxford: Oxford UP.

Heim, Irene & Angelika Kratzer. 1998. *Semantics in generative grammar*. Oxford: Blackwell.

Hendriks, H. 1993. *Studied flexibility: Categories and types in syntax and semantics*: University of Amsterdam dissertation. ICCL Dissertation Series.

Higginbotham, James. 1986. Linguistic theory and Davidson's program in semantics. In E. Lepore (ed.), *The philosophy of Donald Davidson: Perspectives on truth and interpretation*, 29–48. Oxford: Blackwells.

Hodges, Wilfrid. 1998. Compositionality is not the problem. *Logic and Logical Philosophy* 6. 7–33.

Hodges, Wilfrid. 2001. Formal features of compositionality. *Journal of Logic, Language, and Information* 10. 7–28.

Hodges, Wilfrid. 2006. Two doors to open. In Dov Gabbay, Sergei Goncharov & Michael Zakharyaschev (eds.), *Mathematical problems from applied logic I: New logics for the 21st century*, 277–316. New York: Springer.

Hodges, Wilfrid. 2012a. Formalizing the relationship between meaning and syntax. In M. Werning, W. Hinzen & E. Machinery (eds.), *The Oxford handbook of compositionality*, 245–261. Osford: Oxford UP.

Hodges, Wilfrid. 2012b. Requirements on a theory of sentence and word meanings. In R. Schantz (ed.), *Prospects for meaning*, 583–608. New York: de Gruyter.

Husserl, Edmund. 1900. *Logische Untersuchungen,* BAND II. Berlin: Halle. Section references to the translation (of the 1913 2nd Edition) by J.M. Findlay, 1970, as *Logical Investigations, Vol. 2*, Routledge and Kegan Paul, London.

Jacobson, Pauline. 2002. The (dis)organization of the grammar: 25 years. *Linguistics and Philosophy* 25. 601–626.

Jacobson, Pauline. 2012. Direct compositionality. In M. Werning, W. Hintzen & E. Machery (eds.), *The Oxford handbook of compositionality*, 109–128. Oxford: Oxford UP.

Janssen, Theo. 1986. *Foundations and applications of Montague grammar*. Amsterdam: Mathematisch Centrum, Univ. Amsterdam. Especially Chapter 1, "The Principle of Compositionality of Meaning".

Janssen, Theo. 1997. Compositionality. In J. van Benthem & A. ter Meulen (eds.), *Handbook of logic and linguistics*, 417–473. Amsterdam: North Holland.

Janssen, T.M.V. 2001. Frege, contextually and compositonality. *Journal of Logic, Language, and Information* 10. 115–136.

Janssen, T.M.V. 2012. Compositionality: Its historic context. In M. Werning, W. Hintzen & E. Machery (eds.), *The Oxford handbook of compositionality*, 19–46. Oxford: Oxford UP.

Johnson, Michael. 2014. Compositionality. Internet Encyclopedia of Philosophy. http://www.iep.utm.edu/composit.

Kamp, Hans & Barbara H. Partee. 1995. Prototype theory and compositionality. *Cognition* 57. 129–191.

Kaplan, David. 1968/1972. What is Russell's theory of descriptions? In D. Pears (ed.), *The philosophy of Bertrand Russell*, 227–244. Oxford: Oxford UP. Originally presented at a American Philosophical Association (Pacific) meeting, 1968. Published in Pears in 1972.

Kaplan, David. 1989. Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry & H. Wettstein (eds.), *Themes from Kaplan*, 481–566. Oxford: Oxford UP.

Katz, Jerrold & Jerry Fodor. 1963. The structure of a semantic theory. *Language* 39. 170–210.

Kazmi, Ali & Francis Jeffry Pelletier. 1998. Is compositionality formally vacuous? *Linguistics and Philosophy* 21. 629–633.

Lahav, Ran. 1989. Against compositionality: The case of adjectives. *Philosophical Studies* 57. 261–279.

Langacker, Ron. 1987. *Foundations of cognitive grammar. Vol. 1: Theoretical prerequisites*. Stanford: Stanford UP. Chapter 12: Composition.

Larson, Richard & Gabriel Segal. 1995. *Knowledge of meaning: An introduction to semantic theory*. Cambridge: MIT Press.

Montague, Richard. 1970. Universal grammar. *Theoria* 36. 373–398. Reprinted in R. Thomason (ed.) *Formal Philosophy*, (New Haven: Yale Univ. Press), 1974, pp. 222-246.

Nunberg, Geoffrey, Ivan Sag & Thomas Wasaw. 1994. Idioms. *Language* 70. 491–538.

Osgood, Charles, George Suci & Percy Tannenbaum. 1957. *The measurement of meaning*. Urbana, IL: University of Illinois Press.

Pagin, Peter. 1997. Is compositionality compatible with holism? *Mind and Language* 12. 11–33.

Pagin, Peter. 2006. Meaning holism. In Ernest Lepore & Barry Smith (eds.), *The Oxford handbook of philosophy of language*, 213–231. Oxford: Oxford UP.

Pagin, Peter. 2009. Compositionality, understanding, and proofs. *Mind* 118. 713–737.

Pagin, Peter & Dag Westerståhl. 2010a. Compositionality I: Definitions and variants. *Philosophy Compass* 5. 250–264.

Pagin, Peter & Dag Westerståhl. 2010b. Compositionality II: Arguments and problems. *Philosophy Compass* 5. 265–282.

Pagin, Peter & Dag Westerståhl. 2010. Pure quotation and general compositionality. *Linguistics and Philosophy* 33. 381–415.

Parsons, Terence. 1968. A semantics for English. Unpublished manuscript widely circulated in 1968–1973.

Parsons, Terence. 1970. Some problems concerning the logic of grammatical modifiers. *Synthèse* 21. 320–334.

Partee, Barbara. 1984. Compositionality. In F. Landman & F. Veltman (eds.), *Varieties of formal semantics*, 281–312. Dordrecht: Foris. Reprinted in B.H. Partee (2004) *Compositionality in Formal Semantics. Selected papers by Barbara H. Partee*, (Blackwell: Malden, MA), pp. 153–181.

Pelletier, Francis Jeffry. 1994a. On an argument against semantic compositionality. In D. Prawitz & D. Westerståhl (eds.), *Logic and philosophy of science in Uppsala*, 599–610. Dordrecht: Kluwer Academic Publishers.

Pelletier, Francis Jeffry. 1994b. The principle of semantic compositionality. *Topoi* 13. 11–24. Reprinted, with additions, in S. Davis & B. Gillon (2004) *Semantics: A Reader*, Oxford UP, Oxford; pp. 133–156.

Pelletier, Francis Jeffry. 2000. Semantic compositionality: Free algebras and the argument from ambiguity. In M. Feller, S. Kaufmann & M. Pauly (eds.), *Formalizing the dynamics of information*, 207–218. Stanford: CSLI Press.

Pelletier, Francis Jeffry. 2001. Did Frege believe Frege's principle? *Journal of Logic, Language, and Information* 10. 87–114.

Pelletier, Francis Jeffry. 2012. Holism and compositionality. In M. Werning, W. Hinzen & E. Machery (eds.), *The Oxford handbook of compositionality*, 149–174. Oxford: Oxford UP.

Pelletier, Francis Jeffry. forthcoming. Compositionality and concepts from a linguistics and philosophy of language perspective. In J. Hampton & Y. Winter (eds.), *Concepts and compositionality*, Berlin: Springer. Forthcoming 2016.

Potts, Chris. 2007. The dimensions of quotation. In C. Barker & P. Jacobson (eds.), *Direct compositionality*, 405–431. Oxford: Oxford UP.

Predelli, Stefano. 2005. *Contexts: Meaning, truth, and the use of language*. Oxford: Oxford UP.

Predelli, Stefano. 2008. The demonstrative theory of quotation. *Linguistics and Philosophy* 31. 555–572.

Quine, Willard Van Orman. 1940. *Mathematical logic*. Cambridge, MA: Harvard UP.

Quine, Willard Van Orman. 1960. *Word and object*. Cambridge MA: MIT Press.

Recanati, Françis. 2001. Open quotation revisited. *Philosophical Perspectives* 22. 443–471.

Recanati, Françis. 2008. Open quotation. *Mind* 110. 637–687.
Recanati, François. 2012. Compositionality, flexibility, and context dependence. In Markus Werning, Wolfram Hinzen & Edouard Machery (eds.), *The Oxford handbook of compositionality*, 175–191. Oxford: Oxford UP.
Reimer, Marga. 2002. Do adjectives conform to compositionality? *Philosophical Perspectives* 16. 183–.
Robbins, Philip. 2002. How to blunt the sword of compositionality. *Noûs* 36. 313–334.
Russell, Bertrand. 1918. Lectures on logical atomism: Lectures I-II. *The Monist* 28. 495–527. Lectures III–VIII in Volume 29, pp. 32–63, 190–222, 345–380. Also reprinted in R. Marsh (ed) , *Logic and Knowledge: Essays 1901–1950* (Allen & Unwin: London) 1956, and in D. Pears (ed) *Bertrand Russell: The Philosophy of Logical Atomism* (Taylor & Francis e-Library) 2009.
Sag, Ivan & Thomas Warsaw. 1999. *Syntactic theory: A formal introduction*. Stanford: CSLI Press.
Sainsbury, Mark. 2001. Two ways to smoke a cigarette. *Ratio* 14. 386–406.
Saka, P. 1998. Quotation and the use-mention distinction. *Mind* 107. 113–136.
Schank, Roger C. 1972. Conceptual dependency: A theory of natural language understanding. *Cognitive Psychology* 3. 532–631.
Schiffer, Stephen. 1972. *Meaning*. Oxford: Clarendon Press.
Schlick, Moritz. 1918. *General theory of knowledge*. New York: Springer-Verlag. Translated by A.E. Blumberg; published 1974.
Searle, John. 1969. *Speech acts: An essay in philosophy of language*. Cambridge: Cambridge UP.
Searle, John. 1980. The background of meaning. In J. Searle, F. Kiefer & M. Bierwisch (eds.), *Speech act theory and pragmatics*, 221–232. New York: Free Press.
Shan, Chung-Chieh. 2010. The character of quotation. *Linguistics and Philosophy* 33. 417–443.
Smuts, Jan. 1926. *Holism and evolution*. London: MacMillan.
Szabó, Zoltan. 2012. The case for compositionality. In M. Werning, W. Hinzen & E. Machery (eds.), *The Oxford handbook of compositionality*, 64–80. Oxford: Oxford UP.
Tarski, Alfred. 1936. The concept of truth in formalized languages. In J. Corcoran (ed.), *Logic, semantics, metamathematics*, 152–278. Indianapolis: Hackett Pub. Co.
Tarski, Alfred. 1944. The semantic conception of truth. *Phenomenology and Phenomenological Research* 4. 341–376.
Travis, Charles. 1997. Pragmatics. In Bob Hale & Crispin Wright (eds.), *A companion to the philosophy of language*, 87–107. Oxford: Blackwell.

Washington, Cory. 1992. The identity theory of quotation. *Journal of Philosophy* 89. 582–605.

Werning, Markus. 2005. Right and wrong reasons for compositionality. In E. Machinery, M. Werning & G. Schurz (eds.), *The compositionality of meaning and context: Vol. I: Foundational issues*, 285–309. Frankfurt: Ontos.

Westerståhl, Dag. 1998. On mathematical proofs of the vacuity of compositionality. *Linguistics and Philosophy* 21. 635–643.

Westerståhl, Dag. 1999. On the compositionality of idioms: An abstract approach. In D. Barker-Plummer, D. Beaver, J. van Benthem & P. Scotto di Luzio (eds.), *Words, proofs, and diagrams*, 241–271. Stanford: CSLI Press.

Westerståhl, Dag. 2004. On the compositional extension problem. *Linguistics and Philosophy* 33. 549–582.

Westerståhl, Dag. 2012. Compositionality in Kaplan-style semantics. In Markus Werning, Wolfram Hinzen & Edouard Machery (eds.), *The Oxford handbook of compositionality*, 192–219. Oxford: Oxford UP.

Wierzbicka, Anna. 1996. *Semantics: Primes and universals*. NY: Oxford UP.

Wittgenstein, Ludwig. 1921. *Tractatus logico-philosophicus*. London: Routledge & Kegan Paul. Translated by D.F. Pears & B.F. McGuiness; published 1961.

Zadrozny, Wlodek. 1994. From compositional to systematic semantics. *Linguistics and Philosophy* 17. 329–342.

Zimmermann, Thomas Ede. 2012. Compositionality problems and how to solve them. In M. Werning, W. Hinzen & E. Machery (eds.), *The Oxford handbook of compositionality*, 81–106. Oxford: Oxford UP.