Peter Pagin

Stockholm University

Department of Philosophy

106 91 Stockholm, Sweden

phone: +46-8-162813, fax: +46-8-152226

email: peter.pagin@philosophy.su.se

Dag Westerståhl

University of Gothenburg

Department of Philosophy, Linguistics and Theory of Science

Box 100, S-405 30 Gothenburg, Sweden

phone: +46 - 31-786 4573, fax: +46 - 31-786 4853

6. Compositionality

Keywords: compositionality, semantics, syntax

6. Compositionality

Abstract

This article is concerned with the principle of compositionality, i.e. the principle that the meaning of a complex expression is a function of the meanings of its parts and its mode of composition. After a brief historical background, a formal algebraic framework for syntax and semantics is presented. In this framework, both syntactic operations and semantic functions are (normally) partial. Using the framework, the basic idea of compositionality is given a precise statement, and several variants, both weaker and stronger, as well as related properties, are distinguished. Several arguments for compositionality are discussed, and the standard arguments are found inconclusive. Also, several arguments against compositionality, and for the claim that it is a trivial property, are discussed, and are found to be flawed. Finally, a number of real or apparent problems for compositionality are considered, and some solutions are proposed.

1. Background

Compositionality is a property that a language may have and may lack, namely the property that the meaning of any complex expression is *determined* by the meanings of its parts and the way they are put together. The language can be natural or formal, but it has to be interpreted. That is, meanings, or more generally, *semantic values* of some sort must be assigned to linguistic expressions, and compositionality concerns precisely the distribution of these values.

Particular semantic analyses that are in fact compositional were given already in antiquity, but apparently without any corresponding general conception. For instance, in *Sophist*, chapters 24–26, Plato discusses subject-predicate sentences, and suggests (pretty much) that such a sentence is true [false] if the predicate (verb) attributes to what the subject (noun) signifies things that are [are not]. Notions that approximate the modern concept of compositionality did emerge in medieval times. In the Indian tradition, in the 4th or 5th century CE, Śabara says that

> The meaning of a sentence is based on the meaning of the words.

and this is proposed as the right interpretation of a sūtra by Jaimini from sometime 3rd-6th century BCE (cf. Houben 1997, 75–76). The first to propose a general principle of this nature in the Western tradition seems to have been Peter Abelard (2008, 3.00.8) in the first half of the 12th century, saying that

> Just as a sentence materially consists in a noun and a verb, so too the understanding of it is put together from the understandings of its parts. (Translation by and information from Peter King 2007, 8.)

Abelard's principle directly concerns only subject-predicate sentences, it concerns the understanding process rather than meaning itself, and he is unspecific about the nature of the putting-together operation. The high scholastic conception is different in all three respects. In early middle 14th century John Buridan (1998, 2.3, Soph. 2 Thesis 5, QM 5.14, fol. 23vb) states what has become known as the *additive principle*:

> The signification of a complex expression is the sum of the signification of its non-logical terms.(Translation by and information from Peter King 2001, 4).

The additive principle, with or without the restriction to non-logical terms, appears to have become standard during the late middle ages (for instance, in 1500, Peter of Ailly refers to the common view that it 'belongs to the [very] notion of an expression that every expression has parts each one of which, when separated, signifies something of what is signified by the whole'; 1980, 30). The medieval theorists apparently did not possess the general concept of a function, and instead proposed a particular function, that of summing (collecting). Mere collecting is inadequate, however, since the sentences *All A's are B's* and *All B's are A's* have the same parts, hence the same collection of part-meanings and hence by the additive principle have the same meaning.

With the development of mathematics and concern with its foundations came a renewed interest in semantics. Gottlob Frege is generally taken to be the first person to have formulated explicitly the notion of compositionality and to claim that it is an essential feature of human language (although some writers have doubted that Frege really expressed, or really believed in, compositionality; e.g. Pelletier 2001 and Janssen 2001). In "Über Sinn und Bedeutung", 1892, he writes:

> Let us assume for the time being that the sentence has a reference. If we now replace one word of the sentence by another having the same reference, this can have no bearing upon the reference of the sentence (Frege 1892, 62).

This is (a special case of) the *substitution version* of the idea of semantic values being determined; if you replace parts by others with the same value, the value of the whole doesn't change. Note that the values here are *Bedeutun-*

gen (referents), such as truth values (for sentences) and individual objects (for individual-denoting terms).

Both the substitution version and the function version (see below) were explicitly stated by Rudolf Carnap in (1956) (for both extension and intension), and collectively labeled 'Frege's Principle'.

The term 'compositional', was introduced by Hilary Putnam (a former student of Carnap's) in (Putnam 1975a), p. 77, read in Oxford in 1960 but not published until in the collection (Putnam 1975b). Putnam says "[...] the concept of a compositional mapping should be so defined that the range of a complex sentence should depend on the ranges of sentences of the kinds occurring in the 'derivational history' of the complex sentence.". The first use of the term in print seems to be due to Jerry Fodor (a former student of Putnam's) and Jerrold Katz (1964), to characterize meaning and understanding in a similar sense.

Today, compositionality is a key notion in linguistics, philosophy of language, logic, and computer science, but there are divergent views about its exact formulation, methodological status, and empirical significance. To begin to clarify some of these views we need a framework for talking about compositionality that is sufficiently general to be independent of particular theories of syntax or semantics and yet allows us to capture the core idea behind compositionality.

2. Grammars and semantics

The function version and the substitution version of compositionality are two sides of the same coin: that the meaning (value) of a compound expression is a *function* of certain other things (other meanings (values) and a 'mode of composition'). To formulate these versions, two things are needed: a set of *structured expressions* and a *semantics* for them.

Structure is readily taken as *algebraic* structure, so that the set E of linguistic expressions is a domain over which certain syntactic operations or rules are defined, and moreover E is *generated* by these operations from a subset A of *atoms* (e.g. words). In the literature there are essentially two ways of fleshing out this idea. One, which originates with Montague (see 1974a), takes as primitive the fact that linguistic expressions are grouped into *categories* or *sorts*, so that a syntactic rule comes with a specification of the sorts of each argument as well as of the value. This use of *many-sorted algebra* as an abstract linguistic framework is described in Janssen (1986) and Hendriks (2001). The other approach, first made precise in Hodges (2001), is one-sorted but uses *partial algebras* instead, so that rather than requiring the arguments of an operation to be of certain sorts, the operation is simply *undefined* for unwanted arguments. (A many-sorted algebra can in a straightforward way be turned into a one-sorted partial one (but not always vice versa), and under a natural condition the sorts can be recovered in the partial algebra (see Westerståhl 2004 for further details and discussion. Some theorists combine partiality with primitive sorts; for example, Keenan & Stabler 2004 and Kracht 2007.) The partial approach is in a sense simpler and more general than the many-sorted one, and we follow it here.

Thus, let a *grammar*

$$\mathbf{E} = (E, A, \Sigma)$$

be a partial algebra, where E and A are as above and Σ is a set of partial functions over E of finite arity which generate all expressions in E from A. To illustrate, the familiar rules

$$\text{NP} \longrightarrow \text{Det N} \qquad \text{(NP-rule)}$$

$$\text{S} \longrightarrow \text{NP VP} \qquad \text{(S-rule)}$$

correspond to binary partial functions, say $\alpha, \beta \in \Sigma$, such that, if *most*, *dog*, and *bark* are atoms in A, one derives as usual the sentence *Most dogs bark* in E, by first applying α to *most* and *dog*, and then applying β to the result of that and *bark*. These functions are necessarily partial; for example, β is undefined whenever its second argument is *dog*.

It may happen that one and the same expression can be generated in more than one way, i.e. the grammar may allow *structural ambiguity*. So it is not really the expressions in E but rather their *derivation histories*, or *analysis trees*, that should be assigned semantic values. These derivation histories can be represented as *terms* in a (partial) *term algebra* corresponding to $\mathbf{E}$, and a *valuation function* is then defined from terms to surface expressions (usually finite strings of symbols). However, to save space we shall ignore this complication here, and formulate our definitions as if semantic values were assigned directly to expressions. More precisely, the simplifying assumption is that each expression is generated in a *unique* way from the atoms by the rules. One consequence is that the notion of a *subexpression* is well-defined: the subexpressions of t are t itself and all expressions used in the generation of t from atoms (it is fairly straightforward to lift the uniqueness assumption, and reformulate the definitions given here so that they apply to terms in the term algebra instead; see e.g. Westerståhl 2004 for details).

The second thing needed to talk about compositionality is a *semantics* for $\mathbf{E}$. We take this simply to be a function μ from a subset of E to some set M of semantic values ('meanings'). (In the term algebra case, μ takes grammatical terms as arguments. Alternatively, one may take *disambiguated* expressions such as phrase structure markings by means of labeled brackets. Yet another option is to have an extra syntactic level, like Logical Form, as the semantic function domain. The choice between such alternatives is largely irrelevant from the point of view of compositionality.)

The semantic function μ is also allowed to be partial. For example, it may represent our partial understanding of some language, or our attempts at a semantics for a fragment of a language. Further, even a complete semantics will be partial if one wants to maintain a distinction between *meaningfulness* (being in the domain of μ) and *grammaticality* (being derivable by the grammar rules).

No assumption is made about meanings. What matters for the abstract notion of compositionality is not meanings as such, but *synonymy*, i.e. the partial equivalence relation on E defined by:

$$u \equiv_\mu t \text{ iff } \mu(u), \mu(t) \text{ are both defined and } \mu(u) = \mu(t).$$

(We use s, t, u, with or without subscripts, for arbitrary members of E.)

3. Variants and properties of compositionality

3.1. Basic compositionality

Both the function version and the substitution version of compositionality can now be easily formulated, given a grammar $\mathbf{E}$ and a semantics μ as above.

Funct(μ) For every rule $\alpha \in \Sigma$ there is a meaning operation r_α such that if $\alpha(u_1, \ldots, u_n)$ is meaningful, then

$$\mu(\alpha(u_1, \ldots, u_n)) = r_\alpha(\mu(u_1), \ldots, \mu(u_n)).$$

Note that Funct(μ) presupposes the *Domain Principle* (DP): subexpressions of meaningful expressions are also meaningful. The substitution version of compositionality is given by

Subst($\equiv_\mu$) If $s[u_1, \ldots, u_n]$ and $s[t_1, \ldots, t_n]$ are both meaningful expressions, and if $u_i \equiv_\mu t_i$ for $1 \leq i \leq n$, then $s[u_1, \ldots, u_n] \equiv_\mu s[t_1, \ldots, t_n]$.

The notation $s[u_1, \ldots, u_n]$ indicates that s contains (not necessarily immediate) *disjoint* occurrences of subexpressions among $u_1, \ldots, u_n$, and $s[t_1, \ldots, t_n]$ results from replacing each u_i by t_i. (Restricted to immediate subexpressions Subst($\equiv_\mu$) says that $\equiv_\mu$ is a partial *congruence relation*:

> If $\alpha(u_1, \ldots, u_n)$ and $\alpha(t_1, \ldots, t_n)$ are both meaningful and $u_i \equiv_\mu t_i$ for $1 \leq i \leq n$, then $\alpha(u_1, \ldots, u_n) \equiv_\mu \alpha(t_1, \ldots, t_n)$.

Under DP, this is equivalent to the unrestricted version.)

Subst($\equiv_\mu$) does not presuppose DP, and one can easily think of semantics for which DP fails. However, a first observation is:

(1) Under DP, Funct(μ) and Subst($\equiv_\mu$) are equivalent.

That Rule(μ) implies Subst($\equiv_\mu$) is obvious when Subst($\equiv_\mu$) is restricted to immediate subexpressions, and otherwise proved by induction over the generation complexity of expressions. In the other direction, the operations r_α must be found. For $m_1, \ldots, m_n \in M$, let $r_\alpha(m_1, \ldots, m_n) = \mu(\alpha(u_1, \ldots, u_n))$ if there are expressions u_i such that $\mu(u_i) = m_i$, $1 \leq i \leq n$, and $\mu(\alpha(u_1, \ldots, u_n))$ is defined. Otherwise, $r_\alpha(m, \ldots, m_n)$ can be undefined (or arbitrary). This is enough, as long as we can be certain that the definition is independent of the choice of the u_i, but that is precisely what Subst($\equiv_\mu$) says.

The requirements of basic compositionality are in some respects not so strong, as can be seen from the following observations:

(2) If μ gives the same meaning to all expressions, then Funct(μ) holds.

(3) If μ gives different meanings to all expressions, then $\mathrm{Funct}(\mu)$ holds.

(2) is of course trivial. For (3), consider $\mathrm{Subst}(\equiv_\mu)$ and observe that if no two expressions have the same meaning, then $u_i \equiv_\mu t_i$ entails $u_i = t_i$, so $\mathrm{Subst}(\equiv_\mu)$, and therefore $\mathrm{Funct}(\mu)$, holds trivially.

3.2. Recursive semantics

The function version of compositional semantics is given by recursion over syntax, but that does not imply that the meaning operations are defined by recursion over *meaning*, in which case we have *recursive semantics*. Standard semantic theories are typically both recursive and compositional, but the two notions are mutually independent. In the recursive case we have:

$\mathrm{Rec}(\mu)$ There is a function b and for every $\alpha \in \Sigma$ an operation r_α such that for every meaningful expression s,

$$\mu(s) = \begin{cases} b(s) & \text{if } s \text{ is atomic} \\ r_\alpha(\mu(u_1), \ldots, \mu(u_n), u_1, \ldots, u_n) & \text{if } s = \alpha(u_1, \ldots, u_n) \end{cases}$$

For μ to be recursive, the basic function b and the meaning composition operation r_α must themselves be recursive, but this is not required in the function version of compositionality. In the other direction, the presence of the expressions $u_1, \ldots, u_n$ themselves as arguments to r_α has the effect that the compositional substitution laws need not hold (cf. Janssen 1997).

If we drop the recursiveness requirement on b and r_α, $\mathrm{Rec}(\mu)$ becomes vacuous. This is because $r_\alpha(m_1, \ldots, m_n, u_1, \ldots, u_n)$ can simply be *defined* to be $\mu(\alpha(u_1, \ldots, u_n))$ whenever $m_i = \mu(u_i)$ for all i and $\alpha(u_1, \ldots, u_n)$ is meaningful (and undefined otherwise). Since inter-substitution of synonymous but distinct expressions changes at least one argument of r_α, no counterexample is possible.

3.3. Weaker versions

Basic (first-level) compositionality takes the meaning of a complex expresion to be determined by the meanings of the *immediate* subexpressions and the top-level syntactic operation. We get a weaker version — *second-level* compositionality — if we require only that the operations of the *two* highest levels, together with the meanings of expressions at the second level, determine the meaning of the whole complex expression. A possible example comes from constructions with quantified noun phrases where the meanings of both the determiner and the restricting noun – i.e. two levels below the head of the construction in question – are needed for semantic composition, a situation that may occur with possessives and some reciprocals. In Peters & Westerståhl (2006, ch. 7) and in Westerståhl (2008) it is argued that, in general, the corresponding semantics is second-level but not (first-level) compositional.

Third-level compositionality is defined analogously, and is weaker still. In the extreme case we have *bottom-level*, or *weak* functional compositionality, if the meaning of the complex term is determined only by the meanings of its atomic constituents and the entire syntactic construction (i.e. the derived operation that is extracted from a complex expression by knocking out the atomic constituents). A function version of this becomes somewhat cumbersome (but see Hodges 2001, sect. 5), whereas the substitution version becomes simply:

AtSubst($\equiv_\mu$) Just like Subst($\equiv_\mu$) except that the u_i and t_i are all atomic.

Although weak compositionality is not completely trivial (a language could lack the property), it does not serve the language users very well: the meaning operation r_α that corresponds to a complex syntactic operation α cannot be predicted from its build-up out of simpler syntactic operations and their corresponding meaning operations. Hence, there will be infinitely many complex syntactic operations whose semantic significance must be learned one by one.

It may be noted here that terminology concerning compositionality is somewhat fluctuating. David Dowty (2007) calls (an approximate version of) weak functional compositionality *Frege's Principle*, and refers to Funct(μ) as *homomorphism compositionality*, or *strictly local compositionality*, or *context-free semantics*. In Larson & Segal (1995), this is called *strong compositionality*. The labels *second-level compositionality*, *third-level*, etc. are not standard in the literature but seem appropriate.

3.4. Stronger versions

We get stronger versions of compositionality by enlarging the domain of the semantic function, or by placing additional restrictions on meaningfulness or on meaning composition operations. An example of the first is Zoltan Szabo's (2000) idea that the same meaning operations define semantic functions in all possible human languages, not just for all sentences in each language taken by itself. That is, whenever two languages have the same syntactic operation, they also associate the same meaning operation with it.

An example of the second option is what Wilfrid Hodges has called the *Husserl property* (going back to ideas in Husserl 1900):

(Huss) Synonymous expressions belong to the same (semantic) category.

Here the notion of category is defined in terms of substitution; say that $u \sim_\mu t$ if, for every s in E, $s[u] \in dom(\mu)$ iff $s[t] \in dom(\mu)$. So (Huss) says that synonymous terms can be inter-substituted without loss of meaningfulness. This is often a reasonable requirement (though Hodges 2001 mentions some putative counter-examples). (Huss) also has the consequence that Subst($\equiv_\mu$) can be simplified to $\text{Subst}_1(\equiv_\mu)$, which only deals with replacing *one* subexpression by another. Then one can replace n subexpressions by applying $\mathrm{Subst}_1(\equiv_\mu)$ n times; (Huss) guarantees that all the 'intermediate' expressions are meaningful.

An example of the third kind is that of requiring the meaning composition operations to be computable. To make this more precise we need to impose more order on the meaning domain, viewing meanings too as given by an algebra $\mathbf{M} = (M, B, \Omega)$, where $B \subseteq M$ is a finite set of *basic meanings*, Ω is a finite set of elementary operations from n-tuples of meanings to meanings, and M is generated from B by means of the operations in Ω. This allows the definition of meaning operations by recursion over M. The semantic function μ is then defined simultaneously by recursion over syntax and by recursion over the meaning domain. Assuming that the elementary meaning operations are computable in a sense relevant to cognition, the semantic function itself is computable.

A further step in this direction is to require that the meaning operations be *easy* to compute, thereby reducing or minimizing the complexity of semantic interpretation. For instance, meaning operations that are either elementary or else formed from elementary operations by function composition and function application would be of this kind (cf. Pagin 2010 for work in this direction).

Another strengthening, also introduced in Hodges (2001), concerns Frege's so-called Context Principle. A famous but cryptic saying by Frege (1884, x) is: "Never ask for the meaning of a word in isolation, but only in the context of a sentence". This principle has been much discussed in the literature (for example, Dummett 1973, Dummett 1981, Janssen 2001, Pelletier 2001), and sometimes taken to *conflict* with compositionality. However, if not seen as saying that words somehow lose their meaning in isolation, it can be taken as a *constraint* on meanings, in the form of what we might call the Contribution Principle:

(CP) The meaning of an expression is the contribution it makes to the meanings of complex expressions of which it is a part.

This is vague, but Hodges notes that it can be made precise with an additional requirement on the synonymy $\equiv_\mu$. Assume (Huss), and consider:

InvSubst$_\exists(\equiv_\mu)$ If $u \not\equiv_\mu t$, there is an expression s such that either exactly one of $s[u]$ and $s[t]$ is meaningful, or both are and $s[u] \not\equiv_\mu s[t]$.

So if two expressions of the same category are such that no complex expression of which the first is a part changes meaning when the first is replaced by the second, they are synonymous. That is, if they make the *same* contribution to *all* such complex expressions, their meanings cannot be distinguished. This can be taken as one half of (CP), and compositionality in the form of Subst$_1(\equiv_\mu)$ as the other.

Remark: Hodges' main application of these notions is to what has become known as the *extension problem*: given a partial compositional semantics μ, under what circumstances can μ be extended to a larger fragment of the language? Here (CP) can be used as a requirement, so that the meaning of a new word w, say, must respect the (old) meanings of complex expressions of which w is a part. This is especially suited to situations when all new items are parts of expressions that already have meanings (cofinality). Hodges defines a corresponding notion of *fregean extension* of μ, and shows that in the situation just mentioned, and given that μ satisfies (Huss), a *unique* fregean extension always exists. Another version of the extension problem is solved in Westerståhl (2004). An abstract account of compositional extension issues is given in Fernando (2005). *End of remark*

We can take a step further in this direction by requiring that replacement of expressions by expressions with *different* meanings *always* changes meaning:

InvSubst$_\forall(\equiv_\mu)$ If for some i, $0 \leq i \leq n$, $u_i \not\equiv_\mu t_i$, then for every expression s, either exactly one of $s[u_1, \ldots, u_n]$ and $s[t_1, \ldots, t_n]$ are meaningful, or both are and $s[u_1, \ldots, u_n] \not\equiv_\mu s[t_1, \ldots, t_n]$.

This disallows synonymy between complex expressions transformable into each other by substitution of constituents at least some of which are non-synonymous, but it does allow synonymous expressions with different structure. Carnap's principle of synonymy as *intensional isomorphism* forbids this, too. With the concept of *intension* from possible-worlds semantics it can be stated as

(RC) $t \equiv_\mu u$ iff

- i) t, u are atomic and co-intensional, or
- ii) for some α, $t = \alpha(t_1, \ldots, t_n)$, $u = \alpha(u_1, \ldots, u_n)$, and $t_i \equiv_\mu u_i$, $1 \leq i \leq n$

(RC) entails both Subst($\equiv_\mu$) and InvSubst$_\forall$($\equiv_\mu$), but is very restrictive. It disallows synonymy between *brother* and *male sibling* as well as between *John loves Susan* and *Susan is loved by John*, and allows different expressions to be synonymous only if they differ at most in being transformed from each other by substitution of synonymous atomic expressions.

(RC) seems too strong. We get an intermediate requirement as follows. First, define *μ-congruence*, $\simeq_\mu$ in the following way:

($\simeq_\mu$) $t \simeq_\mu u$ iff

- i) t or u is atomic, $t \equiv_\mu u$, and neither is a constituent of the other, *or*
- ii) $t = \alpha(t_1, \ldots, t_n)$, $u = \beta(u_1, \ldots, u_n)$, $t_i \simeq u_i$, $1 \leq i \leq n$, and for all $s_1, \ldots, s_n$, $\alpha(s_1, \ldots, s_n) \equiv_\mu \beta(s_1, \ldots, s_n)$, if either is defined.

Then require synonymous expressions to be congruent:

(Cong) If $t \equiv_\mu u$, then $t \simeq_\mu u$.

By (Cong), synonymous expressions cannot differ much syntactically, but they may differ in the two crucial respects forbidden by (RC). (Cong) does not hold for natural language if logically equivalent sentences are taken as synonymous. That it holds otherwise remains a conjecture (but see Johnson 2006).

It follows from (Cong) that meanings are (or can be represented as) structured entities: entities uniquely determined by how they are built, i.e. entities from which constituents can be extracted. We then have projection operations:

(Rev) For every meaning operation $r : E^n \longrightarrow E$ there are projection operations $s_{r,i}$ such that $s_{r,i}(r(m_1, \ldots, m_n)) = m_i$.

Together with the fact that the operations r_i are meaning operations for a compositional semantic function μ, (Rev) has semantic consequences, the main one being a kind of *inverse* functional compositionality:

InvFunct(μ) The syntactic expression of a complex meaning m is determined, up to μ-congruence, by the composition of m and the syntactic expressions of its parts.

For the philosophical significance of inverse compositionality, see sections 4.6 and 5.2 below. For ($\simeq_\mu$), (Cong), InvFunct(μ), and a proof that (Rev) is a consequence of (Cong) (really of the equivalent statement that the meaning algebra is a *free algebra*), see Pagin (2003a). (Rev) seems to be what Jerry Fodor understands by 'reverse compositionality' in e.g. Fodor (2000, 371).

3.5. Direct and indirect compositionality

Pauline Jacobson (2002) distinguishes between *direct* and *indirect* compositionality, as well as between *strong direct* and *weak direct* compositionality. This concerns how the analysis tree of an expression maps onto the expression itself, an issue we have avoided here, for simplicity. Informally, in strong direct compositionality, a complex expression t is built up from sub-expressions (corresponding to subtrees of the analysis tree for t) simply by means of concatenation. In weak direct compositionality, one expression may wrap around another (as *call up* wraps around *him* in *call him up*). In indirect compositionality, there is no such simple correspondence between the composition of analysis trees and elementary operations on strings.

Even under our assumption that each expression has a unique analysis, our notion of compositionality here is indirect in the above sense: syntactic operations may delete strings, reorder strings, make substitutions and add new elements. Strictly speaking, however, the direct/indirect distinction is not a distinction between kinds of semantics, but between kinds of syntax. Still, discussion of it tends to focus on the role of compositionality in linguistics, e.g. whether to let the choice of syntactic theory be guided by compositionality (cf. Dowty 2007 and Kracht 2007. For discussion of the general significance of the distinction, see Barker & Jacobson 2007).

3.6. Compositionality for "interpreted languages"

Some linguists, among them Jacobson, tend to think of grammar rules as applying to *signs*, where a sign is a triple $\langle e, k, m \rangle$ consisting of a string, a syntactic category, and a meaning. This is formalized by Marcus Kracht (see 2003, 2007), who defines an *interpreted language* to be a set L of signs in this sense, and a *grammar* G as a set of partial functions from signs to signs, such that L is generated by the functions in G from a subset of atomic (lexical) signs. Thus, a meaning assignment is built into the language, and grammar rules are taken to apply to meanings as well.

This looks like a potential strengthening of our notion of grammar, but is not really used that way, partly because the grammar is taken to operate independently (though in parallel) at each of the three levels. Let p_1, p_2, and p_3 be the projection functions on triples yielding their first, second, and third elements, respectively. Kracht calls a grammar *compositional* if for each n-ary grammar rule α there are three operations $r_{\alpha,1}$, $r_{\alpha,2}$, and $r_{\alpha,3}$ such that for all signs $\sigma_1, \ldots, \sigma_n$ for which α is defined,

$$\alpha(\sigma_1, \ldots, \sigma_n) =$$
$$\langle r_{\alpha,1}(p_1(\sigma_1), \ldots, p_1(\sigma_n)), r_{\alpha,2}(p_2(\sigma_1), \ldots, p_2(\sigma_n)), r_{\alpha,3}(p_3(\sigma_1), \ldots, p_3(\sigma_n))\rangle$$

and moreover $\alpha(\sigma_1, \ldots, \sigma_n)$ is defined if and only if each $r_{\alpha,i}$ is defined for the corresponding projections.

In a sense, however, this is not really a variant of compositionality but rather another way to organize grammars and semantics. This is indicated by (4) and (5) below, which are not hard to verify.

First, call G *strict* if $\alpha(\sigma_1, \ldots, \sigma_n)$ defined and $p_1(\sigma_i) = p_1(\tau_i)$ for $1 \leq i \leq n$ entails $\alpha(\tau_1, \ldots, \tau_n)$ defined, and similarly for the other projections. All compositional grammars are strict.

(4) Every grammar G in Kracht's sense for an interpreted language L is a grammar (E, A, Σ) in the sense of section 2 (with $E = L$, $A =$ the set of atomic signs in L, and $\Sigma =$ the set of partial functions of G). Provided G is strict, G is compositional (in Kracht's sense) iff each of p_1, p_2, and p_3, seen as assignments of values to signs (so p_3 is the meaning assignment), is compositional (in our sense).

(5) Conversely, if $\mathbf{E} = (E, A, \Sigma)$ is a grammar and μ a semantics for $\mathbf{E}$, let $L = \{\langle u, u, \mu(u)\rangle : u \in dom(\mu)\}$. Define a grammar G for L (with the obvious atomic signs) by letting

$$\alpha(\langle u_1, u_1, \mu(u_1)\rangle, \ldots, \langle u_n, u_n, \mu(u_n)\rangle) = \\ \langle \alpha(u_1, \ldots, u_n), \alpha(u_1, \ldots, u_n), \mu(\alpha(u_1, \ldots, u_n))\rangle$$

whenever $\alpha \in \Sigma$ is defined for $u_1, \ldots, u_n$ and $\alpha(u_1, \ldots, u_n) \in dom(\mu)$ (undefined otherwise). Provided μ is closed under subexpressions and has the Husserl property, μ is compositional iff G is compositional.

3.7. Context dependence

In standard possible-worlds semantics the role of meanings are served by *intensions*: functions from possible worlds to extensions. For instance, the intension of a sentence returns a truth value, when the argument is a world for which the function is defined. Montague (1974b) extended this idea to include not just worlds but arbitrary *indices* i from some set $\mathbf{I}$, as ordered tuples of contextual factors relevant to semantic evaluation. Speaker, time, and place of utterance are typical elements in such indices. The semantic function μ then assigns a meaning $\mu(t)$ to an expression t, which is itself a function such that for an index $i \in \mathbf{I}$, $\mu(t)(i)$ gives an extension as value. Kaplan's (1989) *two-level* version of this first assigns a function (*character*) to t taking certain parts of the index (the *context*, typically including the speaker) to a *content*, which is in turn a function from selected parts of the index to extensions.

In both versions, the usual concept of compositionality straightforwardly applies. The situation gets more complicated when semantic functions themselves take contextual arguments, e.g. if a meaning-in-context for an expression t in context c is given as $\mu(t, c)$. The reason for such a change might be the view that the contextual meanings are contents in their own right, not just extensional fall-outs of the standing, context-independent meaning. But with context as an additional argument we have a new source of variation. The most natural extension of compositionality to this format is given by

C-Funct(μ) For every rule $\alpha \in \Sigma$ there is a meaning operation r_α such that for every context c, if $\alpha(u_1, \ldots, u_n)$ has meaning in c, then $\mu(\alpha(u_1, \ldots, u_n), c) = r_\alpha(\mu(u_1, c), \ldots, \mu(u_n, c))$.

C-Funct(μ) seems like a straightforward extension of compositionality to a contextual semantics, but it can fail in a way non-contextual semantics cannot, by a *context-shift failure*. For we can suppose that although $\mu(u_i, c) = \mu(u_i, c')$, $1 \leq i \leq n$, we still have $\mu(\alpha(u_1, \ldots, u_n), c) \neq \alpha(u_1, \ldots, u_n), c')$. One might see this as a possible result of so-called *unarticulated constituents*. Maybe the meaning of the sentence

(6) It rains

is sensitive to the location of utterance, while none of the *constituents* of that sentence (say, *it* and *rains*) is sensitive to location. Then the contextual meaning of the sentence at a location l is different from the contextual meaning of the sentence at another location l', even though there is no such difference in contextual meaning for any of the parts. This may hold even if substitution of *expressions* is compositional.

There is therefore room for a weaker principle that cannot fail in this way, where the meaning operation *itself* takes a context argument:

C-Funct(μ)$_c$ For every rule $\alpha \in \Sigma$ there is a meaning operation r_α such that for every context c, if $\alpha(u_1, \ldots, u_n)$ has meaning in c, then

$$\mu(\alpha(u_1, \ldots, u_n), c) = r_\alpha(\mu(u_1, c), \ldots, \mu(u_n, c), c).$$

The only difference is the last argument of r_α. Because of this argument, C-Funct$(\mu)_c$ is not sensitive to the counterexample above, and is more similar to non-contextual compositionality in this respect.

This kind of semantic framework is discussed in Pagin (2005); a general format, and properties of the various notions of compositionality that arise, are presented in Westerståhl (2010). For example, it can be shown that (weak) compositionality for contextual meaning entails compositionality for the corresponding standing meaning, but the converse does not hold.

So far, we have dealt with extra-linguistic context, but one can also extend compositional semantics to dependence on *linguistic context*. The semantic value of some particular *occurrence* of an expression may then depend on whether it is an occurrence in, say, an extensional context, or an intensional context, or a hyperintensional context, a quotation context, or yet something else.

A framework for such a semantics needs a set C of *context types*, an initial *null* context type $\theta \in C$ for unembedded occurrences, and a binary function ψ from context types and syntactic operators to context types. If $\alpha(t_1, \ldots, t_n)$ occurs in context type c_i, then $t_1, \ldots, t_n$ will occur in context type $\psi(c_i, \alpha)$. The context type for a particular occurrence t_i^o of an expression t_i in a host expression t is then determined by its immediately embedding operator α_1, *its* immediately embedding operator, and so on until the topmost operator occurrence.

The semantic function μ takes an expression t and a context type c into a semantic value. The only thing that will differ for *linguistic* context from C-Funct$(\mu)_c$ above is that the context of the subexpressions may be different (according to the function ψ) from the context of the containing expression:

LC-Funct$(\mu)_c$ For every $\alpha \in \Sigma$ there is an operation r_α such that for every context c, if $\alpha(u_1, \ldots, u_n)$ has meaning in c, then

$$\mu(\alpha(u_1, \ldots, u_n), c) = r_\alpha(\mu(u_1, c'), \ldots, \mu(u_n, c'), c),$$

where $c' = \psi(c, \alpha)$.

4. Arguments in favor of compositionality

4.1. Learnability

Perhaps the most common argument for compositionality is the argument from *learnability*: A natural language has infinitely many meaningful sentences. It is impossible for a human speaker to learn the meaning of each sentence one by one. Rather, it must be possible for a speaker to learn the entire language by learning the meaning of a finite number of expressions, and a finite number of construction forms. For this to be possible, the language must have a compositional semantics. The argument was to some extent anticipated already in Sanskrit philosophy of language. During the first or second century BC Patañjali writes:

> . . . Bṛhaspati addressed Indra during a thousand divine years going over the grammatical expressions by speaking each particular word, and still he did not attain the end. . . . But then how are grammatical expressions understood? Some work containing general and particular rules has to be composed . . . (Cf. Staal 1969, 501–502. Thanks to Brendan Gillon for the reference.)

A modern classical passage plausibly interpreted along these lines is due to Donald Davidson:

> It is conceded by most philosophers of language, and recently by some lin-

> guists, that a satisfactory theory of meaning must give an account of how the meanings of sentences depend upon the meanings of words. Unless such an account could be supplied for a particular language, it is argued, there would be no explaining the fact that we can learn the language: no explaining the fact that, on mastering a finite vocabulary and a finite set of rules, we are prepared to produce and understand any of a potential infinitude of sentences. I do not dispute these vague claims, in which I sense more than a kernel of truth. Instead I want to ask what it is for a theory to give an account of the kind adumbrated (Davidson 1967, 17).

Properly spelled out, the problem is not that of learning the meaning of infinitely many meaningful sentences (given that one has command of a syntax), for if I learn that they all mean that *snow is white*, I have already accomplished the task. Rather, the problem is that there are infinitely many propositions that are each expressed by some sentence in the language (with contextual parameters fixed), and hence infinitely many equivalence classes of synonymous sentences.

Still, as an argument for compositionality, the learnability argument has two main weaknesses. First, the premise that there are infinitely many sentences that have a determinate meaning although they have never been used by any speaker, is a very strong premise, in need of justification. That is, at a given time t_0, it may be that the speaker or speakers employ a semantic function μ defined for infinitely many sentences, or it may be that they employ an alternative function μ_0 which agrees with μ on all sentences that have in fact been used but is simply *undefined* for all that have not been used. On the alternative hypothesis, when using a new sentence s, the speaker or the community gives some meaning to s, thereby extending μ_0 to μ_1, and so on. Phenomenologically, of course, the new sentence seemed to the speakers to come already equipped with meaning, but that was just an illusion. On this alternative hypothesis, there is no infinite semantics to be learned. To argue that there is a learnability problem, we must first justify the premise that we employ an infinite semantic function. This cannot be justified by induction, for we cannot infer from finding sentences meaningful that they were meaningful before we found them, and exactly that would have to be the induction base.

The second weakness is that even with the infinity premise in place, the conclusion of the argument would be that the semantics must be computable, but computability does not entail compositionality, as we have seen.

4.2. Novelty

Closely related to the learnability argument is the argument from *novelty*: speakers are able to understand sentences they have never heard before, which is possible only if the language is compositional.

When the argument is interpreted so that, as in the learnability argument, we need to explain how speakers reliably *track* the semantics, i.e. assign to new sentences the meaning that they independently have, then the argument from novelty shares the two main weaknesses with the learnability argument.

4.3. Productivity

According to the pure argument from *productivity*, we need an explanation of why we are able to *produce* infinitely many meaningful sentences, and compositionality offers the best explanation. Classically, productivity is appealed to by Noam Chomsky as an argument for generative grammar. One of the passages runs

> The most striking aspect of linguistic competence is what we may call the ‘creativity of language’, that is, the speaker’s ability to produce new sentences that are immediately understood by other speakers although they bear no physical resemblance to sentences that are ‘familiar’. The

> fundamental importance of this creative aspect of normal language use has been recognized since the seventeenth century at least, and it was the core of Humboldtian general linguistics (Chomsky 1971, 74).

This passage does not appeal to pure productivity, since it makes an appeal to the understanding by other speakers (cf. Chomsky 1980, 76–78). The pure productivity aspect has been emphasized by Fodor (e.g. 1987, 147–148), i.e. that natural language can *express* an open-ended set of propositions.

However, the pure productivity argument is very weak. On the premise that a human speaker can think indefinitely many propositions, all that is needed is to assign those propositions to sentences. The assignment does not have to be systematic in any way, and all the syntax that is needed for the infinity itself is simple concatenation. Unless the assignment is to meet certain conditions, productivity requires nothing more than the combination of infinitely many propositions and infinitely many expressions.

4.4. Systematicity

A related argument by Fodor (1987, 147–150) is that of systematicity. It can be stated either as a property of speaker understanding or as an expressive property of a language. Fodor tends to favor the former (since he is ultimately concerned with the mental). In the simplest case, Fodor points out that if a speaker understands a sentence of the form tRu, she will also understand the corresponding sentence uRt, and argues that this is best explained by appeal to compositionality.

Formally, the argument is to be generalized to cover the understanding of any new sentence that is formed by recombination of constituents that occur, and construction forms that are used, in sentences already understood. Hence, in this form it reduces to one of three different arguments; either to the argument from *novelty*, or to the *productivity* argument, or finally, to the argument from intersubjectivity (below), and only spells out a bit the already familiar idea of old parts in new combinations.

It might be taken to add an element, for it not only aims at explaining the understanding of new sentences that is in fact manifested, but also predicts what new sentences will be understood. However, Fodor himself points out the problem with this aspect, for if there is a sentence s formed by a recombination that we do *not* find meaningful, we will not take it as a limitation of the systematicity of our understanding, but as revealing that the sentence s is not in fact meaningful, and hence that there is nothing to understand. Hence, we cannot come to any other conclusion than that the systematicity of our understanding is maximal.

The systematicity argument can alternatively be understood as concerning natural language itself, namely as the argument that sentences formed by grammatical recombination are meaningful. It is debatable to what extent this really holds, and sentences (or so-called sentences) like Chomsky's *Colorless green ideas sleep furiously* have been used to argue that not all grammatical sentences are meaningful.

But even if we were to find meaningful all sentences that we find grammatical, this does not in itself show that compositionality, or any kind of systematic semantics, is needed for explaining it. If it is only a matter of assigning some meaning or other, without any further condition, it would be enough that we can think new thoughts and have a disposition to assign them to new sentences.

4.5. Induction on synonymy

We can observe that our synonymy intuitions conform to Subst($\equiv_\mu$). In case after case, we find the result of substitution synonymous with the original expression, if the new part is taken as synonymous with the old. This forms the basis of an *inductive generalization* that such substitutions are always meaning preserving. In contrast to the argument from *novelty*, where the idea of tracking the semantics is central, this induction argument may concern our habits of assigning meaning to, or reading meaning into, new sentences: we tend to do it compositionally.

There is nothing wrong with this argument, as far as it goes, beyond what is in general problematic with induction. It should only be noted that the conclusion is weak. Typically, arguments for compositionality aim at the conclusion that there is a systematic pattern to the assignment of meaning to new sentences, and that the meaning of new sentences can be computed somehow. This is not the case in the *induction* argument, for the conclusion is compatible with the possibility that substitutivity is the *only* systematic feature of the semantics. That is, assignment to meaning of new sentences may be completely random, except for respecting substitutivity. If the substitutivity version of compositionality holds, then (under DP) so does the function version, but the semantic function need not be computable, and need not even be finitely specifiable. So, although the argument may be empirically sound, it does not establish what arguments for compositionality usually aim at.

4.6. Intersubjectivity and communication

The problems with the idea of tracking semantics when interpreting new sentences can be eliminated by bringing in intersubjective agreement in interpretation. For by our common sense standards of judging whether we understand sentences the same way or not, there is overwhelming evidence (e.g. from discussing broadcast news reports) that in an overwhelming proportion of cases, speakers of the same language interpret new sentences *similarly*. This convergence of interpretation, far above chance, does not presuppose that the sentences heard were meaningful before they were used. The phenomenon needs an explanation, and it is reasonable to suppose that the explanation involves the hypothesis that the meaning of the sentences are computable, and so it isn't left to guesswork or mere intuition what the new sentences mean.

The appeal to intersubjectivity disposes of an unjustified presupposition about semantics, but two problems remain. First, when encountering new sentences, these are almost invariably produced by a speaker, and the speaker has intended to convey something by the sentence, but the speaker hasn't *interpreted* the sentence, but fitted it to an antecedent thought. Secondly, we have an argument for computability, but not for compositionality.

The first observation indicates that it is at bottom the success rate of linguistic communication with new sentences that gives us a reason for believing that sentences are systematically mapped on meanings. This was the point of view in Frege's famous passage from the opening of 'Compound Thoughts':

> 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 terrestrial 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 thoughts corresponding to the parts of a sentence, so that the structure of the sentence serves as the image of the structure of the thought. (Frege 1923, 55)

As Frege depicts it here, the speaker is first entertaining a new thought, or proposition, finds a sentence for conveying that proposition to a hearer, and by means of that sentence the hearer comes to entertain the same proposition as the speaker started out with. Frege appeals to semantic structure for explaining how this is possible. He claims that the proposition has a structure that mirrors the structure of the sentence (so that the semantic relation may be an isomorphism), and goes on to claim that without this structural correspondence, communicative success with new propositions would not be possible.

It is natural to interpret Frege as expressing a view that entails that compositionality holds as a consequence of the isomorphism idea. The reason Frege went beyond compositionality (or homomorphism, which does not require a one-one relation) seems to be an intuitive appeal to symmetry: the speaker moves from proposition to sentence, while the hearer moves from sentence to proposition. An isomorphism is a one-one relation, so that each relatum uniquely determines the other.

Because of synonymy, a sentence that expresses a proposition in a particular language is typically not uniquely determined within that language by the proposition expressed. Still, we might want the speaker to be able to *work out* what expression to use, rather searching around for suitable sentences by interpreting candidates one after the other. The inverse functional compositionality principle, InvFunct(μ), of section 3.4, offers such a method. Inverse compositionality is also connected with the idea of structured meanings, or thoughts, while compositionality by itself isn't, and so in this respect Frege is vindicated (these ideas are developed in Pagin 2003a).

4.7. Summing up

Although many share the feeling that there is "more than a kernel of truth" (cf. section 4.1) in the usual arguments for compositionality, some care is required to formulate and evaluate them. One must avoid question-begging presuppositions; for example, if a presupposition is that there is an infinity of propositions, the argument for *that* had better not be that standardly conceived natural or mental languages allow the generation of such an infinite set. Properly understood, the arguments can be seen as inferences to the best explanation, which is a respectable but somewhat problematic methodology. (One usually hasn't really tried many other explanations than the proposed one.)

Another important (and related) point is that virtually all arguments so far only justify the principle that the meaning is computable or recursive, and the principle that up to certain syntactic variation, an expression of a proposition is computable from that proposition. Why should the semantics also be compositional, and possibly inversely compositional? One reason could be that compositional semantics, or at least certain simple forms of compositional semantics, is very *simple*, in the sense that a minimal number of processing steps are needed by the hearer for arriving at a full interpretation (or, for the speaker, a full expression, cf. Pagin 2010), but these issues of complexity need to be further explored.

5. Arguments against compositionality

Arguments against compositionality of natural language can be divided into four main categories:

a) arguments that certain constructions are counterexamples and make the principle false,

b) arguments that compositionality is an empirically vacuous, or alternatively trivially correct, principle,

c) arguments that compositional semantics is not *needed* to account for actual linguistic communication,

d) arguments that actual linguistic communication is not *suited* for compositional semantics.

The first category, that of counterexamples, will be treated in a separate section dealing with a number of problem cases. Here we shall discuss arguments in the last three categories.

5.1. Vacuity and triviality arguments

Vacuity. Some claims about the vacuity of compositionality in the literature are based on mathematical arguments. For example, Zadrozny (1994) shows that for every semantics μ there is a compositional semantics ν such that $\nu(t)(t) = \mu(t)$ for every expression t, and uses this fact to draw a conclusion of that kind. But note that the mathematical fact is itself trivial: let $\nu(t) = \mu$ for each t and the result is immediate from (2) in section 3.1 above (other parts of Zadrozny's results use non-wellfounded sets and are less trivial).

Claims like these tend to have the form: for any semantics μ there is a compositional semantics ν from which μ can be easily recovered. But this too is completely trivial as it stands: if we let $\nu(t) = \langle \mu(t), t \rangle$, ν is 1-1, hence compositional by (3) in section 3.1, and μ is clearly recoverable from ν.

In general, it is not enough that the old semantics can be computed from the new compositional semantics: for the new semantics to have any interest it must *agree* with the old one in some suitable sense. As far as we know there are no mathematical results showing that such a compositional alternative can always be found (see Westerståhl 1998 for further discussion).

Triviality. Paul Horwich (e.g. in 1998) has argued that compositionality is not a substantial property of a semantics, but is *trivially* true. He exemplifies with the sentence *dogs barks*, and says (1998, 156–157) that the meaning property

(7) x means DOGS BARK

consists in the so-called construction property

(8) x results from putting expressions whose meanings are DOG and BARK, in that order, into a schema whose meaning is NS V.

As far as it goes, the compositionality of the resulting semantics is a trivial consequence of Horwich's conception of meaning properties. Horwich's view here is equivalent to Carnap's conception of synonymy as intensional isomorphism. Neither allows that that an expression with different structure or composed from parts with different meanings could be synonymous with an expression that means DOGS BARK. However, for supporting the conclusion that compositionality is trivial, these synonymy conditions must themselves hold trivially, and that is simply not the case.

5.2. Superfluity arguments

Mental processing. Stephen Schiffer (1987) has argued that compositional semantics, and public language semantics altogether, is superfluous in the account of linguistic communication. All that is needed is to account for how the hearer maps his mental representation of an uttered sentence on a mental representation of meaning, and that is a matter of a syntactic transformation, i.e. a translation, rather than interpretation. In Schiffer's example (1987, 192–200), the hearer Harvey is to infer from his belief that

(9) Carmen uttered the sentence 'Some snow is white'

the conclusion that

(10) Carmen said that some snow is white

Schiffer argues that this can be achieved by means of transformations between sentences in Harvey's neural language M. M contains a counterpart α to (9), such that α gets tokened in Harvey's so-called belief box when he has the belief expressed by (9). By an inner mechanism the tokening of α leads to the tokening of β, which is Harvey's M counterpart to (10). For this to be possible for any sentence of the language in question, Harvey needs a translation mechanism that implements a recursive translation function f from sentence representations to meaning representations. Once such a mechanism is in place, we have all we need for the account, according to Schiffer.

The problem with the argument is that the translation function f by itself tells us nothing about communicative success. By itself it just correlates neural sentences of which we know nothing except for their internal correlation. We need another recursive function g that maps the uttered sentence *Some snow is white* on α, and a third recursive function h that maps β on the proposition *that some snow is white*, in order to have a complete account. But then the composed function $h(f(g(\ldots)))$ seems to be a recursive function that maps sentences on meanings (cf. Pagin 2003b).

Pragmatic composition. According to François Recanati (2004), word meanings are put together in a process of *pragmatic* composition. That is, the hearer takes word meanings, syntax and contextual features as his input, and forms the interpretation that best corresponds to them. As a consequence, semantic compositionality is not needed for interpretation to take place.

A main motivation for Recanati's view is the ubiquity of those pragmatic operations that Recanati calls *modulations*, and which intuitively contribute to "what is said", i.e. to communicated content before any conversational implicatures. (Under varying terms and conceptions, these phenomena have been described e.g. by Sperber & Wilson 1992, Bach 1994, Carston 2002 and by Recanati himself.) To take an example from Recanati, in reply to an offer of something to eat, the speaker says

(11) I have had breakfast

thereby saying that she has had breakfast in the morning of *the day of utterance*, which involves a modulation of the more specific kind Recanati calls *free enrichment*, and implicating by means of what she says that she is not hungry. On Recanati's view, communicated contents are always or virtually always pragmatically modulated. Moreover, modulations in general do not operate on a complete semantically derived proposition, but on conceptual constituents. For instance, in (11) it is the property of *having breakfast* that is modulated into *having breakfast this day*, not the proposition as a whole or even the property of *having had breakfast*. Hence, it seems that what the semantics delivers does not feed into the pragmatics.

However, if meanings, i.e. the outputs of the semantic function, are *structured* entities, in the sense specified by (Rev) and InvFunct(μ) of section 3.4, then the last objection is met, for then semantics is able to deliver the arguments to the pragmatic operations, e.g. properties associated with VPs. Moreover, the modulations that are in fact made appear to be controlled by a given semantic structure: as in (11), the modulated part is of the same category and occupies the same slot in the overall structure as the semantically given argument that it replaces. This provides a reason for thinking that modulations operate on a given (syntactically induced) semantic structure, rather than on pragmatically composed material (this line of reasoning is elaborated in Pagin & Pelletier 2007).

5.3. Unsuitability arguments

According to a view that has come to be called *radical contextualism*, truth evaluable content is radically underdetermined by semantics, i.e. by literal meaning. That is, no matter how much a sentence is elaborated, something needs to be added to its semantic content in order to get a proposition that can be evaluated as true or false. Since there will always be indefinitely many different ways of adding, the proposition expressed by means of the sentence will vary from context to context. Well-known proponents of radical contextualism include John Searle (e.g. 1978), Charles Travis (e.g. 1985), and Sperber & Wilson (1992). A characteristic example from Charles Travis (1985, 197) is the sentence

(12) Smith weighs 80 kg

Although it sounds determinate enough at first blush, Travis points out that it can be taken as true or as false in various contexts, depending on what counts as important in those contexts. For example, it can be further interpreted as being true in case Smith weighs

(12′)
a. 80 kg when stripped in the morning
b. 80 kg when dressed normally after lunch
c. 80 kg after being force fed 4 liters of water
d. 80 kg four hours after having ingested powerful diuretic
e. 80 kg after lunch adorned in heavy outer clothing

Although the importance of such examples is not to be denied, their significance for semantics is less clear. It is in the spirit of radical contextualism to minimize the contribution of semantics (literal meaning) for determining expressed content, and thereby the importance of compositionality. However, strictly speaking, the truth or falsity of the compositionality principle for natural language is orthogonal to the truth or falsity of radical contextualism. For whether the meaning of a sentence s is a proposition or not is irrelevant to the question whether that meaning is determined by the meaning of the constituents of s and their mode of composition. The meaning of s may be unimportant but still compositionally determined.

In an even more extreme version, the (semantic) meaning of sentence s in a context c is what the speaker uses s to express in c. In that case meaning itself varies from context to context, and there is no such thing as an invariant literal meaning. Not even the extreme version need be in conflict with compositionality (extended to context dependence), since the substitution properties may hold within each context by itself. Context shift failure, in the sense of section 3.7, may occur, if e.g. word meanings are invariant but the meanings of complex expressions vary between contexts.

It is a further question whether radical contextualism itself, in either version, is a plausible view. It appears that the examples of contextualism can be handled by other methods, e.g. by appeal to pragmatic modulations mentioned in section 5.2 (cf. Pagin & Pelletier 2007), which does allow propositions to be semantically expressed. Hence, the case for radical contextualism is not as strong as it may *prima facie* appear. On top, radical contextualism tends to make a mystery out of communicative success.

6. Problem cases

A number of natural language constructions present apparent problems for compositional semantics. In this concluding section we shall briefly discuss a few of them, and mention some others.

6.1. Belief sentences

Belief sentences offer difficulties for compositional semantics, both real and merely apparent. At first blush, the case for a counterexample against compositionality seems very strong. For in the pair

(13) a. John believes that Fred is a child doctor

b. John believes that Fred is a pediatrician

(13a) may be true and (13b) false, despite the fact that *child doctor* and *pediatrician* are synonymous. If truth value is taken to depend only on meaning and on extra-semantic facts, and the extra-semantic facts as well as the meanings of the parts and the modes of composition are the same between the sentences, then the meaning of the sentences must nonetheless be different, and hence compositionality fails. This conclusion has been drawn by Jeff Pelletier (1994).

What would be the reason for this difference in truth value? When cases such as these come up, the reason is usually that there is some kind of discrepancy in the understanding of the attributee (John) between synonyms. John may e.g. erroneously believe that *pediatrician* only denotes a special kind of child doctors, and so would be disposed to assent to (13a) but dissent from (13b) (cf. Mates 1950 and Burge 1978; Mates took such cases as a reason to be skeptical about synonymy). This is not a decisive reason, however, since it is what the words mean in the sentences, e.g. depending on what *the speaker* means, that is relevant, not what the *attributee* means by those words. The speaker contributes with words and their meanings, and the attributee contributes with his belief contents. If John's belief content matches the meaning of the embedded sentence *Fred is a pediatrician*, then (13b) is true as well, and the problem for compositionality is disposed of.

A problem still arises, however, if belief contents are more fine-grained than sentence meanings, and words in belief contexts are somehow tied to these finer differences in grain. For instance, as a number of authors have suggested, perhaps belief contents are propositions under modes of presentation (see e.g. Burdick 1982, Salmon 1986. Salmon, however, existentially quantifies over modes of presentations, which preserves substitutivity). It may then be that different but synonymous expressions are associated with different modes of presentation.

In our example, John may believe a certain proposition under a mode of presentation associated with *child doctor* but not under any mode of presentation associated with *pediatrician*, and that accounts for the change in truth value.

In that case, however, there is good reason to say that the underlying form of a belief sentence such as (13a) is something like

(14) Bel(John, *the proposition that Fred is a child doctor*, M('Fred is a child doctor'))

where M(-) is a function from a sentence to a mode of presentation or a set of modes of presentation. In this form, the sentence *Fred is a pediatrician* occurs both used and mentioned (quoted), and in its used occurrence, *child doctor* may be replaced by *pediatrician* without change of truth value. Failure of substitutivity is explained by the fact that the surface form fuses a used and a mentioned occurrence. In the underlying form, there is no problem for compositionality, unless caused by quotation.

Of course, this analysis is not obviously the right one, but it is enough to show that the claim that compositionality fails for belief sentences is not so easy to establish.

6.2. Quotation

Often quotation is set aside for special treatment as an exception to ordinary semantics, which is supposed to concern *used* occurrences of expressions rather than *mentioned* ones. Sometimes, this is regarded as cheating, and quotation is proposed as a clear counterexample to compositionality: *brother* and *male sibling* are synonymous, but *'brother'* and *'male sibling'* are not (i.e. the expressions that include the opening and closing quote). Since enclosing an expression in quotes is a syntactic operation, we have a counterexample.

If quoting is a genuine syntactic operation, the syntactic rules include a total unary operator κ such that, for any simple or complex expression t,

(15) $\kappa(t) = \text{‘}t\text{’}$

The semantics of quoted expressions is given simply by

(Q) $\mu(\kappa(t)) = t$

Then, since $t \equiv_\mu u$ does not imply $t = u$, substitution of u for t in $\kappa(t)$ may violate compositionality.

However, such a non-compositional semantics for quotation can be transformed into a compositional one, by adapting Frege's view in (1892) that quotation provides a special context type in which expressions refer to themselves, and using the notion of linguistically context-dependent compositionality from section 3.7 above. We shall not give the details here, only indicate the main steps.

Start with a grammar $\mathbf{E} = (E, A, \Sigma)$ (for a fragment of English, say) and a compositional semantics μ for $\mathbf{E}$. First, extend $\mathbf{E}$ to a grammar containing the quotation operator κ, allowing not only quote-strings of the form *‘John’, ‘likes’*, *“Mary”*, etc., but also things like *John likes ‘Mary’* (meaning that he likes the name), whereas we disallow things like *John ‘likes’ Mary* or *‘John likes’ Mary* as ungrammatical. Let E' be the closure of E under the thus extended operations and κ, and let $\Sigma' = \{\alpha' : \alpha \in \Sigma\} \cup \{\kappa\}$. Then we have a new grammar $\mathbf{E}' = (E', A, \Sigma')$ that incorporates quotation.

Next, extend μ to a semantics μ' for $\mathbf{E}'$, using the semantic composition operations that exist by $\mathrm{Funct}(\mu)$, and letting (Q) above take care of κ.

As indicated, the semantics μ' is *not* compositional: even if Mary is the same person as Sue, *John likes ‘Mary’* doesn't mean the same as *John likes ‘Sue’*.

However, we can extend μ' to a semantics μ'' for $\mathbf{E}'$ which is compositional in the sense of LC-Funct$(\mu)_c$ in section 3.7. In the simplest case, there are two context types: c_u, the *use* context type, which is the default type (the null context), and the quotation context type c_q. The function ψ from context types and operators to context types is given by

$$\psi(c, \beta) = \begin{cases} c & \text{if } \beta \neq \kappa \\ c_q & \text{if } \beta = \kappa \end{cases}$$

for $\beta \in \Sigma'$ and c equal to c_u or c_q. μ'' is obtained by redefining the given composition operations in a fairly straightforward way, so that LC-Funct$(\mu'')_c$ is automatically insured. μ'' then extends μ in the sense that if $t \in E$ is meaningful, $\mu''(t, c_u) = \mu(t)$, and furthermore $\mu''(\kappa(t), c_u) = \mu''(t, c_q) = t$.

So μ'' is compositional in the contextually extended sense. That $t \equiv_\mu u$ holds does not license substitution of u for t in $\kappa(t)$, since t there occurs in a quotation context, and we may have $\mu''(t, c_q) \neq \mu''(u, c_q)$. This approach is further developed in Pagin & Westerståhl (2009).

6.3. Idioms

Idioms are almost universally thought to constitute a problem for compositionality. For example, the VP *kick the bucket* can also mean 'die', but the semantic operation corresponding to the standard syntax of, say, *fetch the bucket*, giving its meaning in terms of the meanings of its immediate constituents *fetch* and *the bucket*, cannot be applied to give the idiomatic meaning of *kick the bucket*.

This is no doubt a problem of some sort, but not necessarily for compositionality. First, that a particular semantic operation fails doesn't mean that no other operation works. Second, note that *kick the bucket* is *ambiguous* between its literal and its idiomatic meaning, but compositionality presupposes non-ambiguous meaning bearers. Unless we take the ambiguity itself to be a problem for compositionality (see the next subsection), we should first find a suitable way to disambiguate the phrase, and only then raise the issue of compositionality.

Such disambiguation may be achieved in various ways. We could treat the whole phrase as a lexical item (an atom), in view of the fact that its meaning has to be learnt separately. Or, given that it does seem to have syntactic structure, we could treat it as formed by a different rule than the usual one. In neither case is it clear that compositionality would be a problem.

To see what idioms really have to do with compositionality, think of the following situation. Given a grammar and a compositional semantics for it, suppose we decide to give some already meaningful phrase a non-standard, idiomatic meaning. Can we then extend the given syntax (in particular, to disambiguate) and semantics in a natural way that preserves compositionality? Note that it is not just a matter of accounting for one particular phrase, but rather for all the phrases in which the idiom may occur. This requires an account of how the syntactic rules apply to the idiom, and to its parts if it has structure, as well as a corresponding semantic account.

But not all idioms behave the same. While the idiomatic *kick the bucket* is fine in *John kicked the bucket yesterday*, or *Everyone kicks the bucket at some point*, it is not good in

(16) The bucket was kicked by John yesterday.

(17) Andrew kicked the bucket a week ago, and two days later, Jane kicked it too.

By contrast, *pull strings* preserves its idiomatic meaning in passive form, and *strings* is available for anaphoric reference with the same meaning:

(18) Strings were pulled to secure Henry his position.

(19) Kim's family pulled some strings on her behalf, but they weren't enough to get her the job.

This suggests that these two idioms should be analyzed differently; indeed the latter kind is called "compositional" in Nunberg, Sag & Wasow (1994) (from which (19) is taken), and is analyzed there using the ordinary syntactic and semantic rules for phrases of this form but introducing instead idiomatic meanings of its parts (*pull* and *string*), whereas *kick the bucket* is called "non-compositional".

In principle, nothings prevents a semantics that deals differently with the two kinds of idioms from being compositional in our sense. Incorporating idioms in syntax and semantics is an interesting task. For example, in addition to explaining the facts noted above one has to prevent *kick the pail* from meaning 'die' even if *bucket* and *pail* are synonymous, and likewise to prevent the idiomatic versions of *pull* and *string* to combine illegitimately with other phrases. For an overview of the semantics of idioms, see Nunberg, Sag & Wasow (1994). Westerståhl (2002) is an abstract discussion of various ways to incorporate idioms while preserving compositionality.

6.4. Ambiguity

Even though the usual formulation of compositionality requires non-ambiguous meaning bearers, the occurrence of ambiguity in language is usually *not* seen as a problem for compositionality. This is because *lexical ambiguity* seems easily dealt with by introducing different lexical items for different meanings of the same word, whereas *structural ambiguity* corresponds to different analyses of the same surface string.

However, it is possible to argue that even though there are clear cases of structural ambiguity in language, as in *Old men and women were released first from the occupied building*, in other cases the additional structure is just an *ad hoc* way to avoid ambiguity. In particular, *scope* ambiguities could be taken to be of this kind. For example, while semanticists since Montague have had no trouble inventing different underlying structures to account for the two readings of

(20) Every critic reviewed four films.

it may be argued that this sentence in fact has just one structural analysis, a simple constituent structure tree, and that meaning should be assigned to that one structure. A consequence is that meaning assignment is no longer functional, but relational, and hence compositionality either fails or is just not applicable. Pelletier (1999) draws precisely this conclusion.

But even if one agrees with such an account of the syntax of (20), abandonment of compositionality is not the only option. One possibility is to give up the idea that the meaning of (20) is a proposition, i.e. something with a truth value (in the actual world), and opt instead for *underspecified meanings* of some kind. Such meanings can be uniquely, and perhaps compositionally, assigned to simple structures like constituent structure trees, and one can suppose that some further process of interpretation of particular utterances leads to one of the possible specifications, depending on various circumstantial facts. This is a form of context-dependence, and we saw in section 3.7 how similar phenomena can be dealt with compositionally. What was there called *standing meaning* is one kind of underspecified meaning, represented as a function from indices to 'ordinary' meanings. In the present case, where several meanings are available, one might try to use the *set* of those meanings instead. A similar but more sophisticated way of dealing with quantifier scope is so-called Cooper storage (see

Cooper 1983). It should be noted, however, that while such strategies restore a functional meaning assignment, the compositionality of the resulting semantics is by no means automatic; it is an issue that has to be addressed anew.

Another option might be to accept that meaning assignment becomes relational and attempt instead to reformulate compositionality for such semantics. Although this line has hardly been tried in the literature, it may be an option worth exploring (For some first attempts in this direction, see Westerståhl 2007).

6.5. Other problems

Other problems than those above, some with proposed solutions, include possessives (cf. Partee 1997; Peters & Westerståhl 2006), the context sensitive use of adjectives (cf. Lahav 1989; Szabó 2001; Reimer 2002), noun-noun compounds (cf. Weiskopf 2007), *unless*+quantifiers (cf. Higginbotham 1986; Pelletier 1994), *any* embeddings (cf. Hintikka 1984), and indicative conditionals (e.g. Lewis 1976).

All in all, it seems that the issue of compositionality in natural language will remain live, important and controversial for a long time to come.

References

Abelard, Peter 2008. *Logica 'ingredientibus'. 3. Commentary on Aristotele's* De Interpretatione. Corpus Christianorum Contiunatio Medievalis, Turnhout: Brepols Publishers.

of Ailly, Peter 1980. *Concepts and Insolubles.* Dordrecht: Reidel. Original published as *Conceptus et Insolubilia*, Paris, 1500.

Bach, Kent 1994. Conversational impliciture. *Mind & Language* 9, 124–62.

Barker, Chris & Pauline Jacobson (eds.) 2007. *Direct Compositionality.* Oxford: Oxford University Press.

Burdick, Howard 1982. A Logical Form for the Propositional Attitudes. *Synthèse* 52, 185–230.

Burge, Tyler 1978. Belief and Synonymy. *Journal of Philosophy* 75, 119–38.

Buridan, John 1998. *Summulae de Dialectica 4. Summulae de Suppositionibus*, volume 10-4 of *Artistarium.* Nijmegen.

Carnap, Rudolf 1956. *Meaning and Necessity.* Chicago, Ill.: The University of Chicago Press., 2 edition.

Carston, Robyn 2002. *Thoughts and Utterances. The Pragmatics of Explicit Communication.* Oxford: Oxford University Press.

Chomsky, Noam 1971. Topics in the Theory of Generative Grammar. In: John Searle (ed.) *Philosophy of Language*, Oxford: Oxford University Press. 71–100.

Chomsky, Noam 1980. *Rules and Representations.* Oxford: Blackwell.

Cooper, Robin 1983. *Quantification and Syntactic Theory.* Dordrecht: D. Reidel.

Davidson, Donald 1967. Truth and Meaning. *Synthese* 17, 304–23. Reprinted in (Davidson 1984), 17-36. Page references to the reprint.

Davidson, Donald 1984. *Inquiries into Truth and Interpretation.* Oxford: Clarendon Press.

Dowty, David 2007. Compositionality as an Empirical Problem. In: (Barker & Jacobson 2007), 23–101.

Dummett, Michael 1973. *Frege. Philosophy of Language.* London: Duckworth.

Dummett, Michael 1981. *The Interpretation of Frege's Philosophy.* London: Duckworth.

Fernando, Tim 2005. Compositionality inductively, co-inductively and contextually. In: Edouard Machery, Markus Werning & Gerhard Schurz (eds.) *The Compositionality of Meaning and Content: Foundational Issues*, Frankfurt/Lancaster: Ontos, volume I. 87–96.

Fodor, Jerry 1987. *Psychosemantics.* Cambridge, Mass.: MIT Press.

Fodor, Jerry 2000. Reply to critics. *Mind & Language* 15, 350–374.

Fodor, Jerry & Jerrold Katz 1964. The Structure of a Semantic Theory. In: Jerry Fodor & Jerrold Katz (eds.) *The Structure of Language*, Englewood Cliffs, NJ.: Prentice-Hall. 479–518.

Frege, Gottlob 1884. *Grundlagen der Arithmetik,* translated by J. Austin as *Foundations of Arithmetic,* 1950. Oxford: Blackwell.

Frege, Gottlob 1892. Über Sinn und Bedeutung. In: Peter Geach & Max Black (eds.) *Translations from the Philosophical Writings of Gottlob Frege,* 1952, Oxford: Blackwell. 56–78.

Frege, Gottlob 1923. Gedankengefüge. In: Peter Geach & R H Stoothof (eds.) *Logical Investigations. Gottlob Frege,* 1977, Oxford: Blackwell. 55–78.

Hendriks, Herman 2001. Compositionality and model-theoretic interpretation. *Journal of Logic, Language and Information* 10, 29–48.

Higginbotham, James 1986. Linguistic Theory and Davidson's Program in Semantics. In: Ernest LePore (ed.) *Linguistic Theory and Davidson's Program in Semantics*, Oxford: Blackwell. 29–48.

Hintikka, Jaakko 1984. A Hundred Years Later: The Rise and Fall of Frege's Influence In Language Theory. *Synthèse* 59, 27–49.

Hinzen, Wolfram, Edouard Machery & Markus Werning (eds.) 2010. *The Oxford Handbook of Compositionality.* Oxford: Oxford University Press.

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

Horwich, Paul 1998. *Meaning*. Oxford: Oxford University Press.

Houben, Jan 1997. The Sanskrit Tradition. In: Wout van Bekkum, Jan Houben, Ineke Sluiter & Kees Versteegh (eds.) *The Sanskrit Tradition*, Amsterdam, Philadelphia: John Benjamin Publishing, volume 82 of *Amsterdam Studies in the Theory and History of Linguistic Science*. 49–145.

Husserl, Edmund 1900. *Logische Untersuchungen II/1,* translated by J. N. Findlay as *Logical Investigations,* 1970. London and Henley: Routledge and Kegan Paul.

Jacobson, Pauline 2002. The (Dis)organisation of the Grammar: 25 years. *Linguistics & Philosophy* 25, 601–26.

Janssen, Theo 1986. *Foundations and Applications of Montague Grammar*. Amsterdam: CWI Tracts 19 and 28.

Janssen, Theo 1997. Compositionality. In: Johan van Benthem & Alice ter Meulen (eds.) *Handbook of Logic and Language*, Amsterdam: Elsevier. 417–73.

Janssen, Theo 2001. Frege, contextuality and compositionality. *Journal of Logic, Language and Information* 10, 115–136.

Johnson, Kent 2006. On the Nature of Reverse Compositionality. *Erkenntnis* 64, 37–60.

Kaplan, David 1989. Demonstratives. In: J Almog, J Perry & H Wettstein (eds.) *Themes from Kaplan*, Oxford: Oxford University Press. 481–563.

Keenan, Edward & Edward Stabler 2004. *Bare Grammar: A Study of Language Invariants*. Stanford: CSLI Publications.

King, Peter 2001. Between Logic and Psychology. John Buridan on Mental Language. Presented at the conference *John Buridan and Beyond*, Copenhagen, September 2001.

King, Peter 2007. Abelard on Mental. *American Catholic Philosophical Quarterly* 81, 169–87.

Kracht, Marcus 2003. *The Mathematics of Language*. Berlin: Mouton de Gruyter.

Kracht, Marcus 2007. The emergence of syntactic structure. *Linguistics & Philosophy* 30, 47–95.

Lahav, Ran 1989. Against Compositionality: The Case of Adjectives. *Philosophical Studies* 57, 261–79.

Larson, Richard & Gabriel Segal 1995. *Knowledge of Meaning. An Introduction to Semantic Theory*. Cambridge, Mass.: MIT Press.

Lewis, David 1976. Probabilities of Conditionals and Conditional Probabilities. *The Philosophical Review* 85, 297–315.

Linsky, Leonard (ed.) 1952. *Semantics and the Philosophy of Language*. Urbana, Ill: University of Illionois Press.

Mates, Benson 1950. Synonymity. *University of California Publications in Philosophy* 25, 201–26. Reprinted in (Linsky 1952).

Montague, Richard 1974a. *Formal Philosophy* (edited and with an introduction by Richmond Thomason). New Haven: Yale University Press.

Montague, Richard 1974b. Pragmatics. In: (Montague 1974a), 95–118. Originally published in Raymond Klibansky (ed.), *Contemporary Philosophy: A Survey*, La Nuova Italia Editrice, Florence, 1986, pp. 102-22.

Nunberg, G., I. Sag & T. Wasow 1994. Idioms. *Language* 70, 491–538.

Pagin, Peter 2003a. Communication and strong compositionality. *Journal of Philosophical Logic* 32, 287–322.

Pagin, Peter 2003b. Schiffer on communication. *Facta Philosophica* 5, 25–48.

Pagin, Peter 2005. Compositionality and context. In: Gerhard Preyer & Georg Peter (eds.) *Contextualism in Philosophy*, Oxford: Oxford University Press. 303–48.

Pagin, Peter 2010. Communication and the Complexity of Semantics. In: (Hinzen, Machery & Werning 2010).

Pagin, Peter & Francis Jeffry Pelletier 2007. Content, context and communication. In: Gerhard Preyer & Georg Peter (eds.) *Context-sensitivity and Semantic Minimalism. New Essays on Semantics and Pragmatics*, Oxford University Press. 25–62.

Pagin, Peter & Dag Westerståhl 2009. Pure Quotation, Compositionality, and the Semantics of Linguistic Context. Submitted.

Partee, Barbara H 1997. The Genitive. A Case Study. Appendix to Janssen 1997. In: Johan van Benthem & Alice ter Meulen (eds.) *The Genitive. A Case Study. Appendix to Janssen 1997.*, New York: Elsevier. 464–70.

Pelletier, Francis Jeffry 1994. The Principle of Semantic Compositionality. *Topoi* 13, 11–24. Expanded version reprinted in S. Davis and B. Gillon (eds.) *Semantics: A Reader* (2004), pp. 133-157.

Pelletier, Francis Jeffry 1999. Semantic compositionality: Free algebras and the argument from ambiguity. In: M. Faller, S. Kaufmann & M. Pauly (eds.) *Proceedings of the Seventh CSLI Workshop on Logic, Language and Computation*, Stanford: CSLI Publications. 207–18.

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

Peters, Stanley & Dag Westerståhl 2006. *Quantifiers in Language and Logic*. Oxford: Oxford University Press.

Putnam, Hilary 1975a. Do True Assertions Correspond to Reality? In: *Mind, Language and Reality. Philosophical Papers vol. 2*, Cambridge: Cambridge University Press. 70–84.

Putnam, Hilary 1975b. *Mind, Language and Reality. Philosophical Papers vol. 2*. Cambridge: Cambridge University Press.

Recanati, François 2004. *Literal Meaning*. Cambridge: Cambridge University Press.

Reimer, Marga 2002. Do adjectives conform to compositionality? *Nous* 16, 183–98.

Salmon, Nathan 1986. *Frege's Puzzle*. Cambridge, Mass.: MIT Press.

Schiffer, Stephen 1987. *Remnants of Meaning*. Cambridge, Mass.: MIT Press.

Searle, John 1978. Literal meaning. *Erkenntnis* 13, 207–24.

Sperber, Dan & Deirdre Wilson 1992. *Relevance. Communication & Cognition*. Oxford: Blackwell, 2 edition.

Staal, J. F. 1969. Sanskrit Philosophy of Language. In: Thomas A. Sebeok (ed.) *Linguistics in South Asia*, The Hague: Mouton, volume 5 of *Current Trends in Linguistics*. 499–531.

Szabó, Zoltán 2000. Compositionality as Supervenience. *Linguistics & Philosophy* 23, 475–505.

Szabó, Zoltan 2001. Adjectives in Context. In: R M Harnich & I Kenesei (eds.) *Adjectives in Context*, Amsterdam: John Benjamin Publishing. 119–46.

Travis, Charles 1985. On What is Strictly Speaking True. *Canadian Journal of Philosophy* 15, 187–229.

Weiskopf, Daniel 2007. Compound Nominals, Context, and Compositionality. *Synthèse* 156, 161–204.

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

Westerståhl, Dag 2002. On the compositionality of idioms; an abstract approach. In: D. Barker-Plummer et al. (ed.) *Words, Proofs, and Diagrams*, Stanford: CSLI Publications. 241–271.

Westerståhl, Dag 2004. On the compositional extension problem. *Journal of Philosophical Logic* 33, 549–582.

Westerståhl, Dag 2007. Remarks on scope ambiguity. In: Elisabeth Ahlsén, Peter Juel Henrichsen, Richard Hirsch, Joakim Nivre, Åsa Abelin, Sven Strömqvist & Shirley Nicholson (eds.) *Communication – Action – Meaning. A Festschrift to Jens Allwood*, Göteborg: Department of Linguistics, University of Gothenburg. 43–55.

Westerståhl, Dag 2008. Decomposing generalized quantifiers. *Review of Symbolic Logic* 1:3, 355–371.

Westerståhl, Dag 2010. Compositionality in Kaplan style semantics. In: (Hinzen, Machery & Werning 2010).

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

Peter Pagin, Stockholm (Sweden)

Dag Westerstål, Gothenburg (Sweden)