Questions about Compositionality*

Dag Westerståhl
Department of Philosophy
Stockholm University

Compositionality is currently discussed mainly in computer science, linguistics, and the philosophy of language. In computer science, it is seen as a desirable design principle. But in linguistics and especially in philosophy it is an *issue*. Most theorists have strong opinions about it. Opinions, however, vary drastically: from the view that compositionality is trivial or empty, or that it is simply false for natural languages, to the idea that it plays an important role in explaining human linguistic competence. This situation is unsatisfactory, and may lead an outside observer to conclude that the debate is hopelessly confused.

I believe there is something in the charge of confusion, but that compositionality is nevertheless an idea that deserves serious consideration, for logical as well as natural languages. In this paper I try to illustrate why, without presupposing extensive background knowledge about the issue.[1]

1 Not a vague concept

Here is Jerry Fodor, a well-known philosopher, on compositionality:

> So not-negotiable is compositionality that I'm not even going to tell you what it is.
> ...
> Nobody knows exactly what compositionality demands, but everybody knows why its demands have to be satisfied. (Fodor, 2001):6

And here is the voice of a renowned linguist, David Dowty:

> I believe that there is not and will not be — any time soon, if ever — a unique precise and "correct" definition of compositionality that

*Thanks to Wilfrid Hodges for helpful remarks, and to Peter Pagin for valuable comments and many years of joint work on compositionality-related issues. Work on this paper was supported by a grant from the Swedish Research Council.

[1] There are by now handbook accounts and journal overviews of compositionality, and I will have to refer to these for many details. A good source is the recent *Handbook of Compositionality* (Hinzen, Machery, & Werning, 2012), which in addition to several useful articles has a bibliography that covers most of what has been published in this area. The surveys (Pagin & Westerståhl, 2010a, 2010b) provide definitions, properties, and overviews of several arguments for and against compositionality.

> all linguists and/or philosophers can agree upon ....
> ...
> I propose that we let the term NATURAL LANGUAGE COMPOSITIONALITY refer to *whatever strategies and principles we discover that natural languages actually do employ to derive the meanings of sentences, on the basis of whatever aspects of syntax and whatever additional information (if any) research shows that they do in fact depend on.* (Dowty, 2007):25,27

Both quotes find compositionality 'non-negotiable', but despair of a definition, either because it would be too complicated, or because theorists would disagree about it. Dowty in effect gives up and suggests using the term in a way that makes natural languages compositional by definition.[2]

An immediate reaction is that this is simply wrong: there are completely precise properties of compositionality of which one can ask whether a natural language has them or not. Or rather, whether the language under a given syntactic and semantic analysis has them or not. And this is of course the catch: questions about compositionality are never completely empirical. They also depend on theory. On the other hand, so do most scientific questions. That doesn't mean they have no answers.

To begin, we should bear in mind the following:

- *Given* a language L with a 'reasonable' syntax that identifies *parts* of complex expressions, and *given* an assignment μ of semantic values ('meanings') to expressions, the question whether μ is compositional *is not vague*.
- Indeed, although there are a few distinct notions of compositionality, *each notion is precise* and allows a definite answer to the question.
- Moreover, these notions are general: they don't depend on *how* the syntax or semantics of L is specified.

These observations (to be made good below) should be ground for some optimism. Of course, the real work lies in specifying the syntax-semantics interface, an enterprise guided by considerations which are empirical as well as theoretical. Indeed, compositionality may be one such consideration. If so, we should avoid mystifying or trivializing it.

2 The guiding intuition

The motivation behind postulating compositionality has always been that it helps explain successful linguistic communication, in particular how speakers

[2] I am being slightly unfair to both Fodor and Dowty: Dowty has interesting things to say in that paper about concrete applications of compositionality, and compositionality is a cornerstone in Fodor's criticism of meaning theories such as the prototype theory. My point is just that they unnecessarily obscure the very idea of compositionality.

apparently effortlessly understand sentences never encountered before. Sentences have both structure and meaning, and the thought is that the meaning somehow can be read off the structure. If you know the meanings of the words, and the rules by which they are put together, and also the meaning building operations corresponding to those rules, then you can figure out the meaning of any correctly construed sentence.

This thought has long historical roots. Classically, meanings are taken to be mental objects: concepts or thoughts in the mind, or at least graspable by the mind.[3] For example, the word "horse" corresponds to the concept or idea HORSE, under which all and only horses fall. The word "every" has a different kind of meaning: it does not itself correspond to a 'clear and distinct idea in the mind', but when combined with e.g. "horse", it yields such an idea (exactly which is often less clear), which in turn can be be combined with, say, the concept RUN, to give the meaning that every horse runs.

To make this precise, you need some mathematics: a notion of structure, applicable to linguistic expressions, and possibly also to meanings. The pioneer is Frege, who applied the notion of a *function*: a concept word like "horse" stands for a function HORSE from objects to *truth values*, "everything" corresponds to a *second-level* function Φ which can take a function F like HORSE as an argument, yielding True whenever F yields True for every object.[4] Details aside, sentences express thoughts, which are structured objects, and the structure of the thought is *reflected* in the structure of the sentence.

Or is it the other way around? Consider a scenario based on a much simplified, but still useful, idea of linguistic communication: A wants to communicate a thought T to B. She finds a sentence S that means T, and utters it. B hears S, and reconstructs T from it. Discussions of compositionality usually focus on the *second* part of this transaction: from linguistic items to meanings. Compositionality is invoked to make this step work, even if B has never heard S before. But it seems that the first part is equally important: A may never have uttered S before, so how does she find it, given T? A natural idea is that that step too is compositional.

So we may need compositionality in both directions. At a suitably abstract level, it is presumably the *same* notion in each case. Here I follow tradition and focus on the direction from syntax to meaning.[5]

While theories of syntax are subject to obvious empirical constraints, it is less clear what the data are for theories of meaning. Modern discussions of compositionality tend to circumvent this problem by making the notion more abstract. What seems to matter for compositionality, one may argue, is not

---

[3]For an exposé of historical expressions of this idea about compositionality, which Hodges calls the *Aristotelian* version, see (Hodges, 2012).

[4]So Φ(HORSE) says that everything is a horse; to say that every horse runs you can either use a conditional, $\Phi(\text{HORSE} \rightarrow \text{RUN})$, or let 'every' correspond to a *binary* second-level function.

[5]Bidirectional compositionality is discussed in (Pagin, 2003), where it is observed that Frege's famous opening paragraph in (Frege, 1923) seems to be about both directions. Fodor hints at similar ideas, using the term 'reverse compositionality', e.g. in (Fodor, 2000). Pagin provides a detailed formal analysis of bidirectionality, in particular of how non-trivial synonyms such as 'brother' and 'male sibling' can be dealt with.

what meanings *are* but the fact that the meanings of complex expressions are *determined* by the meanings of their parts (and the way these parts are syntactically combined). Put differently, replacing parts with the same meaning should not change the meaning of the whole. We arrive at the following two modern formulations of compositionality:

(PC-1) The meaning of a complex expression is determined by the meanings of its immediate parts and the mode of composition.

(PC-2) Appropriately replacing (not necessarily immediate) parts of a complex expression with synonymous expressions preserves meaning.

(As to the role of immediacy, see below.) Note that there is no longer any requirement that meanings be mental objects, or objects which themselves can have parts. Indeed, there is no requirement at all on meanings, except that a notion of *sameness of meaning* (synonymy) is available.

3 Structured expressions

To get started, we need a notion of syntax general enough to cover most common forms of grammar. In fact, very little is required: a notion of *structured expression* with identifiable *constituents*. I will consider two similar but distinct ways to proceed, both due to Wilfrid Hodges.[6]

3.1 Syntactic algebras

Systematic attempts to represent natural language syntax in algebraic terms go back at least to (Montague, 1974 (1970)), where, conforming to linguistic practice, expressions are assigned primitive *categories*, in effect making syntactic algebras *many-sorted*. (Hodges, 2001) uses *partial* algebras instead, a simpler approach in the present context. Moreover, Hodges provides an abstract representation of the link between constituent structure and surface form. Thus, a *syntactic algebra* is a structure

$$\mathbf{E} = (E, \alpha^{\mathbf{E}})_{\alpha \in \Sigma}$$

where E is the set of *expressions* and each symbol α in the *signature* Σ denotes an n-ary partial function $\alpha^{\mathbf{E}}$ on E (for some $n \geq 0$), to be thought of as a grammar rule. Partiality, rather than category assignment, is used to restrict the domain of rules to appropriate arguments. *Atomic* expressions can be identified with 0-ary functions.

Expressions in E can be structurally ambiguous, and operations on expressions may suppress meaningful information, so on this picture the syntactic

[6]There are other abstract theories of structured objects, notably (Aczel, 1990), whose notion of a *replacement system* generalizes both set-theoretic and syntactic structure. It doesn't seem directly applicable to questions of compositionality, however.

objects of semantic interest are not the expressions themselves but their *derivation histories* ('analysis trees'). These are immediately obtained as the terms in the *term algebra* corresponding to $\mathbf{E}$. The inductive definition of the set GT of well-formed *grammatical terms* (a subset of the set of all terms), respecting the partiality constraints, simultaneously yields a (surjective) homomorphism val from GT to $\mathbf{E}$. For example,

$$\alpha(a, \beta(b))$$

(a, b atoms) is grammatical iff $\alpha^{\mathbf{E}}$ is defined for the arguments $(val(a), val(\beta(b)))$, and then

$$val(\alpha(a, \beta(b))) = \alpha^{\mathbf{E}}(val(a), val(\beta(b))) = \alpha^{\mathbf{E}}(val(a), \beta^{\mathbf{E}}(val(b)))$$

If $val(t) = val(u)$ for complex terms $t \neq u$, the expression $val(t)$ may be structurally ambiguous. Lexical ambiguity can be dealt with by adding new atoms to the term algebra, e.g. $\overline{bank_1}$ and $\overline{bank_2}$, with $val(\overline{bank_1}) = val(\overline{bank_2}) = bank$.

We now get the constituent relation for free: it is simply the *subterm* relation. Moreover, syntactic categories can be recovered. For $X \subseteq GT$, define

(1) $t \sim_X u$ iff for all terms $s[t]$, $s[t] \in X \Leftrightarrow s[u] \in X$

($s[t]$ indicates that t is a *subterm occurrence* in s, and $s[u]$ is the result of replacing that occurrence by u.) Syntactic categories can then be construed as equivalence classes of $\sim_{GT}$; indeed, a familiar way of identifying categories is precisely in terms of preservation of grammaticality under replacement.

This format fits Montague Grammar, various forms of Categorial Grammar, not to mention the syntax of most logical languages. It also fits the idea of *direct compositionality* of (Jacobson, 2002) and (Barker & Jacobson, 2007). One aspect of 'directness' consists in restrictions on the functions in Σ (e.g. that only concatenation of strings is allowed), and hence on the mapping val. But the main point is that the semantics runs 'in tandem' with the syntax, which means that *val exists*. In grammars using notions of Movement and Logical Form (LF) (see (Heim & Kratzer, 1998) for a textbook example), there is no such mapping. Meanings are (usually compositionally) assigned to LFs, but the rules for constructing LFs have no semantic counterpart; in particular, there need be no homomorphic connection to surface form.[7]

3.2 Constituent structures

Hodges' recent notion of a *constituent structure* (see (Hodges, 2011, 2012)) distills the bare essentials needed for talking about compositionality, and in particular for his notion of a semantics based on Frege's Context Principle (section

[7]The syntactic algebra format also applies, *mutatis mutandis*, to the currently popular idea of formulating grammar rules as applying to *triples* consisting of a string, a syntactic category, and a meaning; see (Kracht, 2003, 2007) for a formal account. So the meaning assignment is built into the grammar rules, but in practice it can be teased apart, and one can usually go between the two formats in a straightforward way—(Pagin & Westerståhl, 2010a), sect. 3.6 has more details.

7 below). Formally, a constituent structure $(\mathbb{E}, \mathbb{F})$ is quite similar to a syntactic algebra: $\mathbb{E}$ is a set of objects called expressions, and $\mathbb{F}$ is a set of partial functions on $\mathbb{E}$. But the intuition is different: think of the elements of $\mathbb{F}$ as *frames* (which is what they are called), obtained from expressions by deleting some parts, leaving *argument places* that can be filled with other expressions, i.e. those in the domain of the frame. For example, from the sentence

(2) Henry knows some students.

you can get various frames, such as

(3) x knows some students

(4) x knows D students

(5) x knows Q

(6) Henry R some A

By definition, $\mathbb{F}$ is closed under composition, substitution, and contains unit frame 1 (a total identity function on $\mathbb{E}$), but no empty frame (function with empty domain).[8] Thus, syntactic term algebras are a special case, with $\mathbb{E}$ as the set of grammatical terms, and $\mathbb{F}$ as the set of *polynomially definable* partial functions on $\mathbb{E}$, i.e., those definable precisely by leaving out subterm occurrences (replacing them with variables) of grammatical terms.

e is said to be a (*proper*) *constituent* of f iff ($e \neq f$ and) there is a frame F such that $f = F(\ldots, e, \ldots)$. In fact (using (NS) in note 8), F can be assumed to be 1-ary. The relation $\sim_X$ now becomes:

$$e \sim_X f \;\text{ iff for each 1-ary } G \in \mathbb{F},\ G(e) \in X \text{ iff } G(f) \in X$$

Constituent structures start from a quite concrete idea of syntactic structure. But the formal requirements are minimal. For example, there is no guarantee that the 'proper constituent' relation is *transitive*. Transitivity follows if the relation is *wellfounded*, a natural enough assumption, but not part of the definition.[9]

[8] More precisely, *Nonempty Composition* is the following:

(NC) If $F(x_1, \ldots, x_n), G(y_1, \ldots, y_m) \in \mathbb{F}$, and $F(e_1, \ldots, e_{i-1}, G(f_1, \ldots, f_m), e_{i+1}, \ldots, e_n) \in \mathbb{E}$, then

$$F(x_1, \ldots, x_{i-1}, G(y_1, \ldots, y_m), x_{i+1}, \ldots, x_n) \in \mathbb{F}$$

And *Nonempty Substitution* is

(NS) If $F(e_1, \ldots, e_n) \in \mathbb{E}$, then $F(x_1, \ldots, x_{i-1}, e_i, x_{i+1}, \ldots, x_n) \in \mathbb{F}$.

[9] The definition allows the existence of (let us call them) *2-loops*: distinct expressions e, f and frames F, G such that $f = F(e)$ and $e = G(f)$. Then e is a proper constituent of f, which is a proper constituent of e, but no expression is a proper constituent of itself, so transitivity fails. Clearly, wellfoundedness precludes 2-loops (or n-loops for any n). It is not hard to show that the 'proper constituent' relation is transitive if and only if there are no 2-loops.

2-loops are in principle allowed in syntactic algebras $\mathbf{E} = (E, \alpha^{\mathbf{E}})_{\alpha \in \Sigma}$ as well (though they would never appear with standard grammar rules), but grammatical terms are always wellfounded.

Since there are normally several ways to turn a given expression into a frame, we will often have

(7) $F(e_1, \ldots, e_n) = G(f_1, \ldots, f_m)$

for distinct n, m, e_i, f_j, F, G. But if different expressions are inserted into the *same* frame, the idea of a frame seems to require that the results be different. Thus, call a frame F *rigid* iff

(8) $F(e_1, \ldots, e_n) = F(f_1, \ldots, f_n)$ implies $e_i = f_i$ for $1 \leq i \leq n$,

i.e. if it is an *injective* function. This looks like another reasonable requirement on constituent structures (which is satisfied in the special case of term algebras), but again it is not needed in Hodges' account.

4 Meanings

Once the wellformed structured expressions have been identified, we can simply let a *semantics* be any assignment μ of values ('meanings') to these. The semantics is *partial* if the domain of μ is a proper subset of the set of expressions, otherwise *total*.

With the syntactic algebra approach, the structured expressions are, not the surface expressions but the grammatical terms in GT. For a constituent structure $(\mathbb{E}, \mathbb{F})$ on the other hand, the only candidates are the expressions in $\mathbb{E}$. Thus, even when (7) holds, we have one expression and hence at most one semantic value. This means that structural ambiguity is not accounted for within the frame picture; some kind of disambiguation must be supposed to have taken place already.[10] Indeed, this seems to be the main conceptual difference between the two approaches to constituency.

Each semantics μ has a corresponding *synonymy* relation:

(9) $s \equiv_\mu t$ iff $\mu(s) = \mu(t)$

Here the right-hand side means: $\mu(s)$ and $\mu(t)$ are both defined, and equal. (The letters 's','t' stand for terms in the term algebra, but exactly the same definition gives the relation $e \equiv_\mu f$ for expressions $e, f \in \mathbb{E}$.)

$\equiv_\mu$ is a partial equivalence relation. Conversely, every partial equivalence relation $\equiv$ on the set of structured expressions generates a corresponding *equivalence class semantics*: $\mu_\equiv(t) = [t]_\equiv = \{s \colon s \equiv t\}$ provided $[t]_\equiv \neq \emptyset$, undefined otherwise. One easily shows that the buck stops here: $\equiv_{\mu_\equiv} = \equiv$.

5 Compositionality

Now we get precise versions of (PC-1) and (PC-2), in each of the syntactic settings above.

[10] Compare Montague's notion of a *language* in (Montague, 1974 (1970)), which is a pair of a *disambiguated language* (essentially a free syntactic algebra) and an unspecified disambiguation relation.

Compositionality, functional version

(i) A semantics μ for GT, given by a syntactic algebra $(E, \alpha^{\mathbf{E}})_{\alpha \in \Sigma}$, is *compositional* iff for each $\alpha \in \Sigma$ there is an operation r_α such that whenever $\mu(\alpha(t_1, \ldots, t_n))$ is defined,

$$\mu(\alpha(t_1, \ldots, t_n)) = r_\alpha(\mu(t_1), \ldots, \mu(t_n))$$

(ii) A semantics μ for $\mathbb{E}$, relative to a constituent structure $(\mathbb{E}, \mathbb{F})$, is *compositional* iff for each $F \in \mathbb{F}$ there is an operation s_F such that whenever $\mu(F(e_1, \ldots, e_n))$ is defined,

$$\mu(F(e_1, \ldots, e_n)) = s_F(\mu(e_1), \ldots, \mu(e_n))$$

The idea is the same in both cases: the value of a complex expression is *determined* by the values of its parts and the mode of composition. In the term algebra, we look at the *immediate* constituents. This notion is not in general available in constituent structures, so we need a separate condition for each frame. Thus, if the situation in (7) obtains, we must have

$$s_F(\mu(e_1), \ldots, \mu(e_n)) = s_G(\mu(f_1), \ldots, \mu(f_m))$$

Note that both versions of (PC-1) require that the domain of μ is *closed under constituents*. This is not necessary for (PC-2):

Compositionality, substitution version

(i) A partial equivalence relation $\equiv$ on GT is *compositional* iff for each term $s[t_1, \ldots, t_n]$, if $t_i \equiv u_i$ for $1 \leq i \leq n$, and $s[t_1, \ldots, t_n], s[u_1, \ldots, u_n]$ are both in the domain of $\equiv$, then

$$s[t_1, \ldots, t_n] \equiv s[u_1, \ldots, u_n]$$

(ii) A partial equivalence relation $\equiv$ on $\mathbb{E}$ is *compositional* iff for each expression $F(e_1, \ldots, e_n)$, if $e_i \equiv f_i$ for $1 \leq i \leq n$, and $F(e_1, \ldots, e_n), F(f_1, \ldots, f_n)$ are both in the domain of $\equiv$, then

$$F(e_1, \ldots, e_n) \equiv F(f_1, \ldots, f_n)$$

In (i), $t_1, \ldots, t_n$ are *disjoint* subterm occurrences in the complex term: if two subterm occurrences of a term are not disjoint, one is a subterm of the other. Constituent structures can model expressions with overlapping constituents, which allows a simpler formulation, and makes the second claim of the next fact trivial. The first claim is also straightforward, but requires an argument by induction over the complexity of terms.

Fact 1
If $dom(\mu)$ is closed under constituents then, in the syntactic algebra setting as well as in the constituent structure setting, μ is compositional iff $\equiv_\mu$ is compositional.

This is satisfactory since it shows that, under some assumptions, there is just *one* notion of compositionality. Thus, for any grammar or syntactic theory that satisfies the minimal requirement of having a reasonable notion of constituency, and for any proposed assignment of meanings to its expressions, the question of whether this assignment is compositional or not *has a definite answer*. Moreover, the only way of showing that such an assignment is *not* compositional, is to exhibit a complex expression that changes its meaning when some of its constituents are replaced by synonymous ones (wrt the meaning assignment).

Example 2 (adjective-noun combinations)
We can use this to immediately lay to rest certain arguments against compositionality. The extension of some adjective-noun combinations is the intersection of the extension of the adjective and the extension of the noun, for example, *male cat* or *prime number*. But in other cases it is not; cf. *white wine* or *red hair*. This has been taken to show that the Adj N construction is not compositional.[11] But it shows nothing of the sort. The extension of *white wine* can still be *determined* by the extension of *white and* the extension of *wine*, *and* the Adj N construction, even if it is not always intersection. Nor does the example show that *white* means something else in *white wine* than it means in, say, *white paper*. That might be the case, or not, but it has nothing to do with (failure of) compositionality. To repeat, the *only* way to show that the Adj N construction is non-compositional (wrt extension) would be to find an expression Adj_1N_1 and an adjective Adj_2 with the same extension as Adj_1 (or a noun N_2 with the same extension as N_1) such that Adj_2N_1 (or Adj_1N_2 or Adj_2N_2) is well-formed and differs in extension from Adj_1N_1. Such examples may, or may not, exist, but as far as I know none have been suggested.

This is not to say that there are no variant notions of compositionality. One weaker version requires only that the meaning of the *atomic* constituents (words) of a structured expression, and the structure itself, determines its meaning.[12] A precise formulation is obtained by restricting the t_i, u_i and the e_i, f_i in the substitution version of compositionality to atomic constituents, where, in a constituent structure, an expression is atomic iff it has no proper constituents. The usual criticism is that this is too weak to figure in any explanation of speaker competence, since the speaker would have to learn, as it were, not just the grammar rules in Σ, but each of the infinitely many syntactic structures that they generate. But note that this sort of criticism can be levelled at the whole constituent structure approach: in the function version there is one semantic operation for each frame.

The syntactic algebra approach brings out the *generative* aspect of syntax, and thereby of a compositional semantics; the constituent structure approach doesn't, and isn't intended to. But *wellfounded* constituent structures recover

[11] Arguments of this kind occur in the literature, but I refrain from giving references.

[12] (Dowty, 2007):23 calls this variant—for reasons unclear to me—Frege's Principle. (Larson & Segal, 1995) call it 'compositionality', and use 'strong compositionality' for compositionality as defined here.

the generative element: it is then possible to generate $\mathbb{F}$ from a set of primitive frames, and compositionality for the primitive frames implies full compositionality. (But if $(\mathbb{E}, \mathbb{F})$ is not wellfounded, there need not even be any primitive frames, or any atoms.)

That said, it should be noted that full compositionality is still a very weak requirement. The best way to see this is via the following observation.

Fact 3
If a semantics μ is one-one, it is compositional.

(This follows from Fact 1, since $\equiv_\mu$ is then the identity relation.) The observation should not come as a surprise, but it highlights the fact that the word "determine" in (PC-1) just means 'is a function of': it doesn't mean that one is 'able to figure out' the meaning of complex expressions from the meanings of their parts. For that, one must impose extra requirements, notably that the meaning operations are *computable* in some suitable sense.

No doubt the computability aspect is also part of the intuitive motivation for compositionality. Still, it makes sense to isolate a *core meaning* of 'compositionality', as in the above definitions. It is the requirement expressed by (PC-1) or similar formulations. In the literature, it has been called *local* compositionality, *strong* compositionality, *homomorphism* compositionality, but the idea is the same. True, it is a weak requirement. But weak is not the same as trivial or empty.

6 Triviality

Compositionality has been charged with triviality for both mathematical and philosophical reasons. In the former case, the idea is roughly that any semantics can be made compositional by some trivial manipulations. There is a sense in which this is true. It is just that this fact tells us next to nothing about the unmanipulated semantics. The philosophical charge is rather that compositionality adds nothing to an account of linguistic meaning. I will look at one typical example of each kind.[13]

6.1 Mathematical triviality: Zadrozny

(Zadrozny, 1994) shows that given any semantics μ one can find another semantics μ^* with the same domain such that (a) μ^* is compositional; (b) μ can be recovered from μ^*.[14] In fact, the semantics μ^* is one-one, so its compositionality is indeed trivial (Fact 3). But the claim that a semantics satisfying (a) and (b)

[13] Part of the discussion in this section comes from (Westerståhl, 1998) and (Pagin & Westerståhl, 2010b), where several other examples are examined as well.

[14] He also shows that with a non-wellfounded set theory as metatheory, the only composition operation required for μ^* is function application. This is more interesting, but irrelevant to the issue of the triviality of compositionality.

exists is *itself* trivial: just let, for each $e \in dom(\mu)$,

$$\mu'(e) = (\mu(e), e)$$

Then μ' is compositional (since it is one-one), and μ is easily recovered from μ' ($\mu(e)$ is the first element of the pair $\mu'(e)$). Clearly, this says *nothing at all* about the original semantics μ.

A very different observation is that it has often happened that a proposed semantics μ has been replaced by a compositional semantics μ^*, precisely because μ turned out not to be compositional. Perhaps the first example is Frege's introduction of indirect Sinn and Bedeutung in order to be able to deal (compositionally) with attitude reports. A recent case is Hodges' compositional trump semantics for the Hintikka-Sandu Independence-Friendly Logic, (Hodges, 1997). These semantics are not obtained by trivial manipulations but by a deeper analysis of meaning.

If there is anything in the charge of triviality for mathematical reasons it comes from the observation in Fact 3. When the analysis of meaning is so fine-grained that there are no non-trivial synonymies, compositionality is indeed trivial. To take an extreme example, if the sound of the words themselves, or the associations they conjure up in the mind of the speaker, are taken to be part of the meaning expressed, very few distinct expressions will mean the same. This is not a notion of meaning for which compositionality makes a difference. It doesn't follow that there aren't others for which it does.

6.2 Philosophical triviality: Horwich

In "Deflating compositionality" in (Horwich, 2005), Paul Horwich accepts compositionality but gives it no role whatsoever in explaining the meaning of complex sentences. The idea is that the meanings of words (atoms) and the rules of syntax provide all the information needed:

(a) That x means DOGS BARK consists in x resulting from putting together words whose meanings are DOGS and BARK, in that order, into a schema whose meaning is NS V.

(b) "dogs" means DOGS, "bark" means BARK, and "ns v" means NS V.

(c) "dogs bark" results from putting "dogs" and "bark", in that order, into the schema "ns v".

(d) Hence, "dogs bark" means DOGS BARK.

Horwich's conclusion is that compositionality holds as a direct consequence of what it is for a complex expression to have meaning.

I think the possible attraction of this argument comes from the fact that the example is so simple that any meaning explanation is bound to appear trivial. Looking closer, however, this impression dissolves.

First, one may wonder if the idea is that no other string of words can mean DOGS BARK, and similarly for other sentences. If so, we have trivial compositionality because of a one-one meaning assignment, as just discussed. But that is not the reason offered. Second, the reason this is unclear is that we are not told what the meanings of DOGS or BARK are, and even less about the operation of concatenating two such meanings. Is the notation used a shorthand for a semantic operation of combining the meaning of a bare plural with the meaning of an intransitive verb? Compositionality says that such an operation *exists*. But the order of explanation is the reverse: *after* we have specified such an operation (not done in (a)–(d)), we can *conclude* that compositionality holds.

Third, the example may look trivial but the compositionality claim still has content. It says that other sentences of the *same* form, for example "Cats meow", should be analyzed with the *same* semantic operation. If you find that too trivial, you have an *argument* for compositionality!

Finally, the appearance of triviality fades with more complex sentences:

(10) Everyone knows someone.

It is easy to specify schemas generating (10). It is less trivial to specify corresponding semantic operations that yield the intended meaning (rather, one of the intended meanings) of (10), though, of course, nowadays every semanticist knows ways to do that. To say that the meaning of (10) is EVERYONE KNOWS SOMEONE is completely uninformative until the semantic operations are specified. To say that language requires such operations to exist is to *presuppose* compositionality. But then it looks like an essential trait of language, and anything but trivial. However, it seems more fruitful to regard it as an *hypothesis* about natural language meaning. After all, it is easy to make up non-compositional languages. So it is a substantial hypothesis, to which empirical evidence is relevant. It may look 'deflated' with examples like the one in (a), but it really isn't.

6.3 Triviality: conclusion

Even if there are various uninteresting ways to make a non-compositional semantics compositional, isn't it a significant fact that in so many cases, what looked like non-compositional linguistic constructions have been amenable to a compositional treatment? To evaluate the significance of this, one would have to look at the instances case by case, and there is no space for that here. But, hypothetically, suppose that in each case it was in fact possible to replace the non-compositional semantics by an *improved* semantics which was compositional. That would certainly count as evidence for the *truth* of the compositionality hypothesis. Or, suppose instead that some constructions would resist a compositional treatment. This need not mean we must give up compositionality altogether; it could still be that large *fragments* of natural languages are compositional.

In the second case, compositionality is surely not trivial: it would be false for some parts of language and true for others. What about the first case? For all we have said so far, it could still be that compositionality is trivially true, in the sense that on the ultimately best account of how language works, it plays no significant explanatory role. But we are not at that point yet. In the meantime, it still looks like an hypothesis worth exploring further.

Besides, once we have a well-defined framework in which to talk about compositionality, several related but distinct issues suggest themselves. We look at one in the next section.

7 Hodges and the Context Principle

Frege's second methodological maxim in the introduction to *Grundlagen der Aritmetik* famously reads:

> Nach der Bedeutung der Wörter muß im Satzzusammenhange, nicht in ihrer Vereinzelung gefragt werden. (Frege, 1884):

Frege's application was that the meaning of *number words* is given by the sentences in which they occur, but the general idea seems to be:

(F) The meaning of an expression is the contribution it makes to the meanings of sentences in which it occurs.

Hodges observed that this is in fact a recipe for recovering expression meanings, up to synonymy, from sentence meanings.[15] Let a language L be given as a constituent structure $(\mathbb{E}, \mathbb{F})$ with a semantics μ, where $X = dom(\mu)$. For the next definition, recall sections 3.2 and 4.

Definition 4 (fregean semantics)
For $e, f \in \mathbb{E}$, define

$e \equiv^F_\mu f$ iff
$e \sim_X f$ and for each 1-ary $G \in \mathbb{F}$, if $G(e) \in X$ then $G(e) \equiv_\mu G(f)$

Note that $\equiv^F_\mu$ is a total equivalence relation on $\mathbb{E}$. Let $|e|_\mu$ be the equivalence class of e (alternatively, a chosen label for that class); this is called the *fregean semantics* for L.

[15] See (Hodges, 2001, 2005). Hodges is one of those who have contributed most to our understanding of compositionality, so it is no accident that his name appears so often in this paper. Apart from contributions mentioned here, Hodges resolved the issue of the compositionality of Hintikka's Independence-Friendly (IF) Logic (Hintikka, 1996; Hintikka & Sandu, 1997), by providing it with a compositional semantics (Hodges, 1997), but also showing (Cameron & Hodges, 2001) that no semantics with sets of assignments as values (as for first-order logic) is compositional (see (Galliani, 2011) for a strengthened version of this). His compositional so-called trump semantics sparked off a surge of research on logics where notions of (in)dependence are treated explicitly, notably Dependence Logic (DL) (Väänänen, 2007); see also (Kontinen, Väänänen, & Westerståhl, to appear). He has also contributed significantly to our knowledge of the *history* of the idea of compositionality, especially in Arabic medieval philosophy commenting on Aristotle, but also its modern history with Frege and Tarski. And he has applied his mathematical insights to careful discussion of various linguistic constructions.

Lemma 5 (Hodges' Lifting Lemma)
Suppose $F(e_1,\dots,e_n)$ is a constituent of some expression in X, and $e_i \equiv_\mu^F f_i$ for each i. Then

(a) $F(f_1,\dots,f_n) \in \mathbb{E}$

(b) $F(e_1,\dots,e_n) \equiv_\mu^F F(f_1,\dots,f_n)$

Proof. (outline) The fact that $\mathbb{F}$ is closed under substitution allows us to restrict attention to the case $n = 1$. The assumption about $F(e_1)$, and that $e_1 \sim_X f_1$, together with the fact that $\mathbb{F}$ is closed under composition, yields (a). (b) follows by a similar argument. □

A crucial property of the set of sentences is that it is *cofinal*: every expression is a constituent of some sentence. So if we assume that $X = dom(\mu)$ is cofinal, the Lifting Lemma immediately shows that the fregean semantics is compositional. Thus (Fact 1), for each $F \in \mathbb{F}$ there is an operation h_F such that whenever $F(e_1,\dots,e_n) \in \mathbb{E}$,

$$|F(e_1,\dots,e_n)|_\mu = h_F(|e_1|_\mu,\dots,|e_n|_\mu)$$

How does the fregean semantics relate to the original semantics μ? By Definition 4, we get (since the unit frame belongs to $\mathbb{F}$):

(11) If $e \in X$ and $e \equiv_\mu^F f$, then $e \equiv_\mu f$ (so $f \in X$).

That is, $\equiv_\mu^F$ *refines* $\equiv_\mu$: it may make more meaning distinctions than $\equiv_\mu$ does, but it will never declare synonymous two expressions in X that are not μ-synonymous. However, if μ is already compositional, and satisfies what Hodges calls the *Husserl property*,

(12) if $e \equiv_\mu f$, then $e \sim_X f$

(recall that $X = dom(\mu)$), then it follows that $\equiv_\mu^F$ *coincides* with $\equiv_\mu$ on X.[16] This in fact means that it is possible to choose a label $\nu(e)$ for each $|e|_\mu$ such that ν *extends* μ, i.e. for $e \in X$, $\nu(e) = \mu(e)$. In other words, under these circumstances, the meaning of sentences is unchanged, and the fregean semantics extends the given meaning assignment to (all) other expressions of L. If we in addition assume that the constituent structure of L is wellfounded (section 3.2), Hodges observes (the *Abstract Tarski Theorem*) that the fregean semantics can be presented as recursive definition, with base clauses for atomic expressions, and clauses for complex expressions of the special form

(13) $\nu(F(e_1,\dots,e_n)) = h_F(\nu(e_1),\dots,\nu(e_n))$

[16] In more detail, we have the following fact, which is immediate from the definitions:

Fact 6 (Hodges)
The following are equivalent:

(a) $\equiv_\mu^F$ *coincides with* $\equiv_\mu$ *on* X.

(b) *For all* $e, f \in X$ *and* $F \in \mathbb{F}$, $e \equiv_\mu f$ *and* $F(e) \in X$ *implies* $F(e) \equiv_\mu F(f)$.

These abstract results already have interesting applications to formal languages: Hodges notes that they establish the existence of a Tarski-style truth definition for IF logic (see note 15) as well as for the (closely related) logic with branching quantifiers (i.e. branching of $\forall$ and $\exists$). What do they tell us about natural languages?

First of all that the Context Principle, in the form (F), is indeed viable. But there are some caveats. One is that the fregean semantics is only defined up to synonymy, so it tells us nothing about what suitable fregean values are. Here Hodges is optimistic: in practice it has turned out that natural ways of finding out when two expressions have different fregean values yield natural ways of choosing suitable labels for the equivalence classes. In any case, if our main interest is compositionality, synonymy is enough.

A seemingly more pressing issue is—again—triviality. One might think that, even if the sentence semantics μ is not one-one, it is fine-grained enough that for any two distinct expressions you can find a sentence such that replacing one by the other in it changes its meaning. If so, the fregean semantics is one-one outside X, and thus essentially trivial. But this is another instance of the fact that you need a substantial notion of synonymy for properties like compositionality to make a difference. In this case, not all nuances of meaning should be taken into account; perhaps sameness of truth conditions, or sameness of the expressed proposition (in some suitable sense), is enough. Moreover, as Hodges notes, it makes sense to restrict attention to fragments of languages, deliberately excluding certain constructions. Rather than as a way to avoid complications, this can be seen as abstracting from some details of reality in order to bring out underlying uniformities, a common procedure in the natural sciences.

Still, what are we to make of the fact that the fregean semantics is *always*—provided $X = dom(\mu)$ is cofinal—compositional? Simply, I think, that this is a feature of the most natural way of recovering expression meanings from sentence meanings. It doesn't in itself have empirical content. But the properties of the fregean semantics tell us, to begin with, to direct our attention to the sentence semantics μ. For only when μ is well behaved, in particular, is itself compositional, will the fregean semantics *extend* μ. Only then is it related in a reasonable way to the semantics we started with. And μ shouldn't be compositional for trivial reasons, and it shouldn't make the fregean semantics trivial either.

Furthermore, the fregean semantics may clarify our reflection on intuitive notions of meaning, or rather synonymy. As Hodges says, we have to solve the equation

$$\equiv_\mu^F \ = \ \sim_X \ \cap \ \approx$$

where $\sim_X$ comes from syntax (provided identifying sentences is a syntactic matter) and $\approx$ is an intuitive synonymy relation. The relation $\equiv_\mu^F$ itself is a trivial solution, but finding more reasonable solutions involves real semantic work (Hodges gives several illustrations). These are the real lessons, it seems to me, from Frege's Context Principle.

8 Quotation: a counter-example?

I will not discuss here the many counter-examples to compositionality that have been proposed, and the compositional solutions that have been suggested. But I will look at one case, which is perhaps the clearest of them all: (pure) quotation, i.e. the ability to refer in the language to linguistic expressions (meaningful or not). In a perfectly clear, and in principle familiar, sense, quotation is not compositional. Let us make this a bit more precise.

A language L is, as above, identified with a constituent structure $(\mathbb{E}, \mathbb{F})$ with a distinguished cofinal set $X \subseteq \mathbb{E}$ of (declarative) sentences, and a semantics μ with domain X. We say that L is *interpreted* if each sentence is either true or false, and that μ *respects truth values* if whenever e and f differ in truth value, $\mu(e) \neq \mu(f)$.

I will further say that L *has quotation* if there is a unary frame $Q \in \mathbb{F}$ such that, intuitively, $Q(e)$ is a quote frame of e (e.g. e surrounded by quotation marks) when $e \in X$,[17] and L is able to express elementary syntactic properties of sentences. The details need not be specified, but the point is there are sentences in L, with $Q(e)$ as a constituent, which are true iff, say, e begins with the letter "a", or e consists of five words, etc. Then we have:

(NQ) *Suppose L is an interpreted language that has quotation and whose sentence semantics μ respects truth values. Then, either μ is one-one or it is not compositional.*

For suppose μ is not one-one, i.e. that there are distinct $e, f \in X$ such that $\mu(e) = \mu(f)$. Since they have distinct shapes, some true sentence s in X with $Q(e)$ as a constituent is sensitive to this difference: it becomes false when e is replaced by f. There is a frame $G \in \mathbb{F}$ such that $s = G(e)$. Since μ respects truth values, $\mu(G(e)) \neq \mu(G(f))$. So μ is not compositional. And so $\equiv_\mu^F$ does not coincide with $\equiv_\mu$ on sentences: we have $e \equiv_\mu f$ but $e \not\equiv_\mu^F f$. Indeed, as remarked in the preceding section, the fregean semantics becomes trivial.

This is essentially nothing but the familiar 'opacity' of quotation, but formulated in general terms which reveal the very minimal assumptions needed about L; for example, it doesn't rely on identifying meaning with reference. There are statements in the literature which appear to contradict (NQ), but on a closer look, they don't.[18]

What should we conclude? The strategy of weakening the synonymy $e \equiv_\mu f$ doesn't seem helpful, since respecting truth values looks like a minimum requirement. The remaining alternative is to simply leave out quotation from the

[17] It is enough to assume here that we can quote sentences. In general, of course, one wants to quote arbitrary expressions, perhaps even arbitrary sequences of atomic symbols.

[18] For example, (Potts, 2007) presents an elegant semantics for (not only pure) quotation, which he claims to be compositional. What he in effect does is to give a recursive truth definition whose clauses for complex expressions are not of the form (13) but rather

$$\nu(F(e_1, \ldots, e_n)) = h_F(\nu(e_1), \ldots, \nu(e_n), e_1, \ldots, e_n)$$

Thus, the expressions themselves, as well as their meanings, are arguments of the semantic operations. This is much weaker than (homomorphism) compositionality; see also (Pagin & Westerståhl, 2010a), sect. 3.2.

language. That is certainly possible. On the other hand, quotation, in the pure form of having a means of referring to linguistic items, is such a natural mechanism with such a straightforward semantics. And if we admit this mechanism in the language, compositionality is lost.

But maybe not completely lost. Section 10 will sketch a generalization of compositionality that admits quotation, and certain other recalcitrant linguistic constructions as well. But first I need to say something about compositionality and *context*.

9 Dependence on extra-linguistic context

Context dependence in natural languages is ubiquitous. The clearest case is *indexicals*. Normally one wants to assign a *meaning* to sentences like

(14) I am hungry.

But if this meaning is to have anything to do with *truth conditions*, you need to account for the fact that the truth of (14) varies with the context of utterance. There are basically two ways to proceed. Either you let the meaning assignment μ take expressions *and* contexts as arguments. Or you *curry*, that is, you introduce, in the words of (Lewis, 1980), 'constant and complex semantic values', values which themselves are functions from contexts to ordinary meanings.[19] On the curried approach the notion of compositionality as we have defined it applies. But on the first approach we have this extra argument, requiring a slight reformulation. How slight, and what are the relations between the two approaches? Abstractly, the situation is easy to describe.[20]

As before, the language L has a constituent structure $(\mathbb{E}, \mathbb{F})$ and a semantics μ, but now μ is a function from $\mathbb{E} \times C$ to some set Z of 'ordinary' meanings, where C is a set of *contexts*. For simplicity, I'll assume μ is total. Contexts can be any objects; typical cases are

- $\mu(\forall x\varphi, f) = \mathrm{T}$ iff for all $a \in M$, $\mu(\varphi, f(a/x)) = \mathrm{T}$ (contexts as assignments)
- $\mu(P\varphi, t) = \mathrm{T}$ iff for some $t' < t$, $\mu(\varphi, t') = \mathrm{T}$ (contexts as times)
- $\mu(I, c) = \mathit{speaker}_c$ (contexts as utterance situations)

Currying, we get the 1-ary function

$$\mu_{curr} : \mathbb{E} \longrightarrow [C \longrightarrow Z]$$

($[X \longrightarrow Y]$ is the set of functions from X to Y), defined by

[19] This is the *functional* version. On the *structured* version, meanings are structured objects, possibly with 'holes' that can be filled with e.g. contexts. Everything I say below holds, with slight alterations, for the structured approach as well.

[20] For a details, motivation, proofs, and discussion of the issues raised in this section, see (Pagin, 2005) and (Westerståhl, 2012). Note that the 'meanings' in Z can themselves be functions, say, from possible worlds to truth values.

$\mu_{curr}(e)(c) = \mu(e, c)$

We know what compositionality of μ_{curr} amounts to. For μ, there are two slightly different natural notions (using the functional formulation):

Context-sensitive compositionality

(i) μ is *compositional* iff for each $F \in \mathbb{F}$ there is an operation s_F such that for all $c \in C$,

$$\mu(F(e_1, \ldots, e_n), c) = s_F(\mu(e_1, c), \ldots, \mu(e_n, c))$$

(ii) μ is *weakly compositional* iff for each $F \in \mathbb{F}$ there is an operation s_F such that for all $c \in C$,

$$\mu(F(e_1, \ldots, e_n), c) = s_F(\mu(e_1, c), \ldots, \mu(e_n, c), c)$$

So the only difference is that the context itself is allowed to be an argument of the semantic operations in the weak case. This is actually an important weakening, and allows as compositional several phenomena that are often considered pragmatic rather than semantic. Here is how these notions are related.

Proposition 7
(Contextual) compositionality of μ implies weak (contextual) compositionality of μ, which in turn implies (ordinary) compositionality of μ_{curr}, but none of these implications can in general be reversed.

The first two examples above, with contexts as assignments and as times, respectively, are typical instances of semantics which are *not* (not even weakly) contextually compositional, but where the curried version *is* compositional. The first of these reflects the familiar fact that Tarski's truth definition for first-order logic is compositional if you take *sets of assignments* (not truth values) as semantic values. The third example, on the other hand, with contexts as utterance situations, you typically expect to belong to a (contextually) compositional semantics. The reason is that in the first two cases contexts are *shifted* in the right-hand side of the clause, but this is usually not thought to happen in the third case.

There is much to say about which notion applies to which kind of linguistic construction, but here the points to take home are these: (a) Compositionality makes perfect sense also when meaning is context-dependent (which is the rule rather than the exception in natural languages). (b) But there are (at least) three distinct notions involved, related as in Proposition 7, and in applications one needs to be aware of which one is at stake.

10 General compositionality

Once extra-linguistic context dependence is seen to be compatible with compositionality, there is no reason why linguistic context dependence shouldn't also be. Such dependence can be understood in different ways. One is dependence on other parts of *discourse*, as when an anaphoric pronoun refers back to something introduced earlier by a name or, as in (15), an indefinite description:

(15) A woman entered the room. Only Fred noticed her.

Here I am interested in dependence on *sentential* context, of the kind Frege talks about in the following well-known passage:

> If words are used in the ordinary way, what one intends to speak of is their reference. It can also happen, however, that one wishes to talk about the words themselves or their sense. This happens, for instance, when the words of another are quoted. One's own words then first designate words of the other speaker, and only the latter have their usual reference. We then have signs of signs. In writing, the words are in this case enclosed in quotation marks. Accordingly, a word standing between quotation marks must not be taken to have its ordinary reference. (Frege, 1892):58–9

What Frege says here is that *the type of linguistic context can change the meaning*. Quotation is one such type, sometimes indicated by quotation marks, and in this context, words no longer refer to what they usually refer to, but to themselves. Attitude contexts is of another type (only hinted at in this passage but developed in other parts of (Frege, 1892)); then we use the same words to "talk about ... their sense."

In the syntactic algebra framework (section 3.1), terms are construction trees, so you can identify the (linguistic) context of a term occurrence t *in* a sentence s (or any complex term with t as a subterm) with the unique *path* from the top node to t. Let a *context typing* be a partition of the set of such paths, with the property that the type of each daughter t_i of a node $\alpha(t_1, \ldots, t_n)$ is determined by the type of that node, α, and i. Then we can formulate compositionality with C as the set of context types just as we did weak compositionality for arbitrary C, but with the difference that the meaning of $\alpha(t_1, \ldots, t_n)$ at c is determined by α, c, and the meanings of the t_i *at* c_i, where c_i is the context type determined by c, α, and i.

This version doesn't easily extend to the constituent structure framework (section 3.2), but there is another formulation, equivalent to the one just sketched for syntactic algebras, but applying more generally.[21] In the constituent structure framework, the idea would be to let a semantics be a *set* S of mappings from $\mathbb{E}$ to meanings, together with a *selection function* Ψ, telling which function $\mu_i \in S$ should be applied to e_i when μ applies to $F(e_1, \ldots, e_n)$. Thus, *compositionality of* (S, Ψ) is the property that for each $F \in \mathbb{F}$ and each $\mu \in S$ there is an operation $r_{\mu,F}$ such that when $F(e_1, \ldots, e_n) \in \mathbb{E}$,

$$\mu(F(e_1, \ldots, e_n)) = r_{\mu,F}(\mu_1(e_1), \ldots, \mu_n(e_n)),$$

[21] This formulation is due to Peter Pagin. For full details of these notions of compositionality (in the syntactic algebra setting), their properties, and the application to quotation, see (Pagin & Westerståhl, 2010c).

where $\mu_i = \Psi(\mu, F, i)$. So there is no extra argument to the meaning assignment, but instead there may be more than one meaning assignment function. We call this *general compositionality*. (If S is a unit set we have the ordinary notion.)

The application to quotation is now straightforward: in the simplest version you just need two meaning assignment functions, a *default* function μ_d and a *quotation* function μ_q, and the quote frame Q (section 8) has the property that whatever function is applied to $Q(e)$, μ_q is applied to e. And of course, for all $e \in \mathbb{E}$, $\mu_q(e)$ is (the surface representation of) e itself.

The idea of a semantics that allows switching between different meaning assignments appears quite natural, not only for quotation but for certain other linguistic phenomena as well.[22] Frege had a similar idea for attitude contexts. (Glüer & Pagin, 2006, 2008, 2012) use such a semantics for the modal operators, to deal with rigidity phenomena without treating names or natural kind terms as rigid designators. The point here has just been to show that compositionality, in its general form, is still a viable issue for such semantics.

11 Summing up

The question about compositionality, given a semantics for a set of structured expressions, is not vague. I have illustrated how it is spelled out relative to two (closely related) abstract accounts of syntax, accounts that fit most current syntactic theories. I have also emphasized that the real work lies in the choice of the syntax/semantics interface. Compositionality can be a factor in this choice, provided it is thought to play a role in an account of how language works. Most semanticists believe that it does. There are dissenting voices, but I have not yet seen a convincing purely mathematical or purely philosophical argument that it is trivial or empty. Nor have I seen a proposed counter-example that is not amenable to a compositional treatment—two kinds of context dependence were given as illustrations. And even if counter-examples should exist, it seems beyond doubt that efforts to insure compositionality have lead to exciting developments in semantics and in logic. That is one kind of evidence that compositionality is a good thing. Another is the cluster of related issues that the study of compositionality brings to light, such as the relation between word meaning and sentence meaning, or the Husserl property (section 7). I think we may conclude that in the present state of language research, it would be ill-advised to disregard the issue of compositionality.

References

Aczel, P. (1990). Replacement systems and the axiomatization of situation theory. In R. Cooper, K. Mukai, & J. Perry (Eds.), *Situation Theory and*

[22]According to (Hodges, 2012):249, this is the notion of *taḥrīf* used by the eleventh-century Persian-Arabic writer Ibn Sīnā ('Avicenna'), among others. Hodges is skeptical of its usefulness in semantics (pp. 255–6), but at least the application to quotation (not discussed by him) seems very natural.

its Applications, vol. 1 (p. 3-31). Stanford: CSLI Publications.

Barker, C., & Jacobson, P. (Eds.). (2007). *Direct Compositionality.* Oxford: Oxford University Press.

Cameron, P., & Hodges, W. (2001). Some combinatorics of imperfect information. *Journal of Symbolic Logic*, *66*, 673–684.

Dowty, D. (2007). Compositionality as an empirical problem. In C. Barker & P. Jacobson (Eds.), *Direct Compositionality* (p. 23-101). Oxford: Oxford University Press.

Fodor, J. (2000). Reply to critics. *Mind & Language*, *15*, 350–374.

Fodor, J. (2001). Language, thought, and compositionality. *Mind & Language*, *16*, 1-15.

Frege, G. (1884). *Grundlagen der Arithmetik,* translated by j. austin as *Foundations of Arithmetic,* 1950. Oxford: Blackwell.

Frege, G. (1892). Über sinn und bedeutung. In P. Geach & M. Black (Eds.), *Translations from the philosophical writings of gottlob frege,* 1952 (pp. 56–78). Oxford: Blackwell.

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

Galliani, P. (2011). Sensible semantics of imperfect information. In M. Banerjee & A. Seth (Eds.), *Logic and Its Applications, vol. 6521 of LNCS* (p. 79-89). Berlin/Heidelberg: Springer.

Glüer, K., & Pagin, P. (2006). Proper names and relational modality. *Linguistics & Philosophy*, 507-35.

Glüer, K., & Pagin, P. (2008). Relational modality. *Journal of Logic, Language and Information*, *17*, 307-22.

Glüer, K., & Pagin, P. (2012). General terms and relational modality. *Nous*, *46*, 159-99.

Heim, I., & Kratzer, A. (1998). *Semantics in Generative Grammar.* Oxford: Blackwell.

Hintikka, J. (1996). *The Principles of Mathematics Revisited.* Cambridge: Cambridge University Press.

Hintikka, J., & Sandu, G. (1997). Game-theoretical semantics. In J. van Benthem & A. ter Meulen (Eds.), *Handbook of Logic and Language* (pp. 361–410). Amsterdam: Elsevier.

Hinzen, W., Machery, E., & Werning, M. (Eds.). (2012). *The Oxford Handbook of Compositionality.* Oxford: Oxford University Press.

Hodges, W. (1997). Compositional semantics for a language of imperfect information. *Logic Journal of the IGPL*, *4(5)*, 539–563.

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

Hodges, W. (2005). A context principle. In R. Kahle (Ed.), *Intensionality* (pp. 42–59). Wellesley Mass.: Association for Symbolic Logic and A. K. Peters.

Hodges, W. (2011). From sentence meanings to full semantics. In J. van Benthem et al. (Ed.), *Proof, Computation and Agency* (pp. 261–276). Springer.

Hodges, W. (2012). Formalizing the relationship between meaning and syntax. In W. Hinzen, E. Machery, & M. Werning (Eds.), *The Oxford Handbook of Compositionality* (pp. 245–261). Oxford: Oxford University Press.

Horwich, P. (2005). *Reflections on meaning*. Oxford: Oxford University Press.

Jacobson, P. (2002). The (dis)organisation of the grammar: 25 years. *Linguistics & Philosophy*, *25*, 601-26.

Kontinen, J., Väänänen, J., & Westerståhl, D. (Eds.). (to appear). *Dependence and Independence in Logic*. Special issue of *Studia Logica*.

Kracht, M. (2003). *The Mathematics of Language*. Berlin: Mouton de Gruyter.

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

Larson, R., & Segal, G. (1995). *Knowledge of Meaning. An Introduction to Semantic Theory*. Cambridge, Mass.: MIT Press.

Lewis, D. (1980). Index, context, and content. In S. Kanger & S. Öhman (Eds.), *Philosophy and Grammar*. Dordrecht: D. Reidel. (Reprinted in D. Lewis, *Papers in Philosophical Logic*, 1998, 21–44. Cambridge: Cambridge University Press)

Montague, R. (1974 (1970)). Universal grammar. In *Formal Philosophy. Selected papers by Richard Montague* (pp. 222–246). Yale University Press.

Pagin, P. (2003). Communication and strong compositionality. *Journal of Philosophical Logic*, *32*, 287–322.

Pagin, P. (2005). Compositionality and context. In G. Preyer & G. Peter (Eds.), *Contextualism in Philosophy* (p. 303-48). Oxford: Oxford University Press.

Pagin, P., & Westerståhl, D. (2010a). Compositionality I. Definitions and variants. *Philosophy Compass*, *5(3)*, 250–264.

Pagin, P., & Westerståhl, D. (2010b). Compositionality II. Arguments and problems. *Philosophy Compass*, *5(3)*, 265–282.

Pagin, P., & Westerståhl, D. (2010c). Pure quotation and general compositionality. *Linguistics & Philosophy*, *33(5)*, 381-415.

Potts, C. (2007). The dimensions of quotation. In C. Barker & P. Jacobson (Eds.), *Direct Compositionality* (p. 405-31). Oxford: Oxford University Press.

Väänänen, J. (2007). *Dependence Logic*. Cambridge: Cambridge University Press.

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

Westerståhl, D. (2012). Compositionality in Kaplan style semantics. In W. Hinzen, E. Machery, & M. Werning (Eds.), *The Oxford Handbook of Compositionality* (pp. 192–219). Oxford: Oxford University Press.

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