

Compositionality and Ambiguity*

Dag Westerståhl
Göteborg University

1 The issue

The topic of this paper is the relation between two principles, both widely endorsed in semantics, which may be preliminarily expressed as follows:

- (C) The meaning of a complex expression is determined by the meanings of its (immediate) constituents and the mode of composition.
- (A) Ambiguity, in the sense that expressions have more than one meaning, is a pervasive linguistic phenomenon.

There is an *apparent* conflict between (C) and (A). You may think it is only apparent, and that it is obvious how it should be resolved. This might be the view of many linguists and philosophers of language. But others have felt that the conflict does pose a real problem.

In general, though, one cannot say that the issue has been discussed very much. What I propose to do in this paper is simply to dig a little more into it. I think there are in fact interesting problems involved, and it is worth while to state them. My first aim will be to illustrate and clarify that claim. The second aim is to propose a new way of reconciling the two principles, and explore some of its consequences.

Before we can begin, however, some preliminary points need to be made.

2 Preliminary Methodological Remarks

2.1 Background Assumptions

To make sense of (C) we need, minimally, some function μ , which I will call here a *semantics*, from some collection E of *structured expressions* to some collection M of ‘meanings’.

*I would like to thank Josh Dever, Peter Pagin, Zoltan Szabo, Jouko Väänänen, and the audience at the Semantics Workshop, Rutgers University, May 2001, for helpful comments on earlier versions of this material. Work on this paper was partially supported by a grant from the Swedish Research Council.

In fact this is all we need. To be completely precise, one should specify how structure is generated, presumably from some set of lexical items or atoms A by means of some syntactic rules. A useful and economic algebraic framework for all of this is presented in Hodges [8], and it forms the technical backdrop of much of what follows, though I will avoid technicalities here as far as possible.

Of course, *using* (C) for various explanatory purposes requires further assumptions (cf. below). But there is a point in seeing that (C) in itself presupposes nothing about what expressions look like, or what meanings are, or where μ comes from. The point is generality: (C) can be applied in many different kinds of situations. All further requirements are extra.

As to (A), I need not rehearse the familiar fact that there are many *kinds* of ambiguity to take into account: lexical, structural, referential, and what have you. But concerning its possibility of being at all compatible with (C), it seems we have but two options: Either we think of structure as (completely) disambiguating, or we relax μ from a function to a *relation* R between expressions and meanings. In the latter case, of course, it is no longer so clear what (C) says. Yet the latter course is precisely the one I want to explore, but before I get there we have a little more work to do.

Staying with the function μ for the moment, let us agree to call an expression p (μ)-*meaningful* if it is in the domain of μ (we need not assume that all well-formed expressions are meaningful), and to say that p and q are (μ)-*synonymous*, in symbols $p \equiv_{\mu} q$, if both are meaningful and $\mu(p) = \mu(q)$.

2.2 Empirical or Methodological?

There is a somewhat popular idea that compositionality, i.e., (C), is a merely methodological principle. Some even claim that the empirical emptiness of (C) can be *proved* by mathematical arguments (in particular, Zadrozny [15]). I have written elsewhere about this (see Westerståhl [13]), and will thus confine myself here to a brief remark.

The supposed ‘proofs’ that (C) is empty have the following form: Given *any* semantics μ for E we can define *another* semantics μ' for E which satisfies (C) and from which μ can be recovered. But this is in fact *trivial* (for example, let the new meaning of p be the pair consisting of the old meaning and p itself), and hence *no* interesting consequences follow.

If you have a counterexample to compositionality, the interesting question is if you can adjust your syntactic analysis and/or your meanings in a way that preserves or increases explanatory power and yet satisfies (C). There are several good instances of this in the history of semantics, but none of them are trivial.¹

¹Think of Frege’s notion of indirect reference, or Montague’s compositional analysis of subject-predicate form, or the replacement of truth conditions by input-output conditions effected in modern dynamic semantics (allowing compositional treatment of anaphora also across sentence borders); cf. Groenendijk and Stokhof [5]. Or, the compositional solutions in Pelletier [10] to the problem of interpreting *unless* raised in Higginbotham [6]. (By calling these interesting and non-trivial I am not implying that they are ultimately successful — that is another issue.)

And surely, the issue of whether it can be done or not is *both* methodological *and* empirical (like most issues in most sciences). To see if the linguistic facts fit your theory, some preparatory theoretical distinctions have to be in place. But not any theory will do; there are criteria for that too. This is just common-sense philosophy of science.

So, I maintain, whether (C) holds or not is *not* an empty question, at least not for any apparent reason. And the same holds, *mutatis mutandis*, for (A), by similarly common-sensical arguments.

2.3 Explaining Understanding and Communication

The principles (C) and (A) are interesting if they are true, and useful if they help explain how linguistic understanding and communications works. Someone might go further and claim that any possible language suitable for human understanding and communication would have to satisfy (C) and/or (A). I don't need such a strong claim; it is enough if actual languages, or significant fragments of them, satisfy the principles.

In particular, (C) is invoked in connection with the productivity, or learnability, or systematicity of language. But clearly (C) by itself cannot explain any of this. It just states the *existence* of certain meaning operations. Extra assumptions are needed to connect them to any human activity, for example that we *know* these meaning operations, or that we *follow* them, or compute by means of them. Thus, (C) should be seen at most as a necessary condition on a theory of understanding or communication. But this is enough to make its investigation, and that of its relation to (A), a worthwhile undertaking.

That may be another trivial point. Still, I'd like to illustrate it with two further remarks.

2.3.1 Unique Meanings

Suppose — counterfactually, I'm sure — that no two expressions have the same meaning, i.e., that our semantics μ is a one-one function. Then (C) is trivially true. It is immediate from any precise version of (C) (cf. section 2.3.2) that the required meaning operations then exist. But that would explain nothing. Even if our infinitely many linguistic expressions all had distinct meanings, there could still be a story to tell about how we 'figure out' the meanings of complex expressions. But now the extra requirements on compositionality do all the work; (C) is just a necessary condition, which in this particular — and unlikely — case becomes trivial.

2.3.2 Meanings vs. Synonymy

There are two versions of (C). One associates with each syntactic rule α a corresponding meaning operation r_α . I shall call it the *rule* version. The other, *substitution* version says something about what happens when we substitute synonymous expressions. Here they are:

Rule(μ) For each syntactic rule α there is a (partial) operation r_α such that whenever a complex term $\alpha(p_1, \dots, p_n)$ is meaningful,

$$\mu(\alpha(p_1, \dots, p_n)) = r_\alpha(\mu(p_1), \dots, \mu(p_n)).$$

Subst(μ) If $p_i \equiv_\mu q_i$ for $1 \leq i \leq n$, and if $p = p[p_1, \dots, p_n]$ is a meaningful expression with constituents (not necessarily immediate) p_1, \dots, p_n , such that $p[q_1, \dots, q_n]$ is also meaningful, then

$$p[p_1, \dots, p_n] \equiv_\mu p[q_1, \dots, q_n].$$

It is often reasonable to assume (though it does not hold for all semantics proposed in the literature) that if an expression is meaningful, so are its constituents. The assumption is sometimes called the Domain Rule. Note that the rule version presupposes the Domain Rule, but the substitution version does not, and the latter is thus a more general version of compositionality. But *if* the Domain Rule holds, then the two versions are equivalent; see Hodges [8] and Westerståhl [14].

One reason I mention this here is that the substitution version does not really use the semantics μ , only the synonymy relation \equiv_μ . That is, it abstracts away from *what meanings are*. But clearly an account of understanding or communication cannot abstract away completely from that.² Hence, once again, (C) is only a necessary condition.

The other reason for this excursion was that I needed to mention the rule version and the substitution version of (C) anyway, since for ambiguous semantics it does matter which one we choose, as will be seen presently.

3 Common Attitudes towards the (C)–(A) Issue

As I said, there is an apparent conflict between (C) and (A). In the literature, when this conflict is noted at all, it is usually met with one or more of the following attitudes:

- (a) Deny the conflict by claiming that compositionality only applies after disambiguation.

This can be seen as the position of classical semantics, like Montague Grammar, for example. We know that we *can* disambiguate, since we have logical languages in which all meanings can be expressed. Such languages are compositional by design. But Montague showed that we can use logical forms which are more faithful to the structure of real languages than the original ones, which were

²A point also insisted on in Hodges [7].

designed for mathematics.³ Of course, by (a), any difference in meaning has to correspond to a difference in structure. But why? What is the motivation for this?

If we think semantics should have something to say about human understanding and communication, (a) looks like a substantial claim. It might be false. Even if it is true, it should be backed, I think, by less question-begging arguments than one saying that we only know how to make sense of (C) for disambiguated forms. (In fact, I will suggest a way to make sense of it for ambiguous forms.) Those arguments should say something about syntactic structure, about the evidence for (perhaps various layers of) structure, and about how we process it.

(C) is about how we use structure to get to meaning. It seems that all that classical semantics can say here is that some parsing process computes all logical forms for us, for which the meaning of each is arrived at compositionally, and then we choose one. That might be computationally very inefficient, hence unrealistic. Perhaps there is some more direct way to get from structure to meaning? At the very least, if such questions are admitted, we cannot simply stick to (a) by stipulation.

- (b) Claim that the conflict is resolved if one switches to what I shall call *set meaning*, where the set meaning of an expression p is the set of all ordinary meanings of p .

This appears to be a very common claim (cf. section 4 below). For example, it essentially underlies the treatment of quantifier scope ambiguity in Cooper [2] ('Cooper storage'). As far as I can see, there is really no reason to expect (C) to automatically hold for set meaning in this sense. But this claim requires argument: one such comes in section 7.1.

- (c) Claim that the conflict can be resolved by using *underspecified* meanings.

This looks promising, and the use of underspecified meaning representations is a main feature of modern computational semantics, one of whose avowed aims is to better explain actual language understanding. However, the status of underspecified meaning representations is not exactly clear. Are they merely technical devices, a compact notation for a set of fully specified representations? Or do they have their own 'meanings', for example, something we can *reason* with? In the former case, we seem to be stuck with set meaning, which, I claim, is not an option. In the latter case, it remains to be specified what those meanings are. To my knowledge, there is no agreement on how to do that.

³For example, pioneers of predicate logic like Frege and Russell were proud to claim — correctly — that in this logic there is no subject-predicate form, and that the adherence to subject-predicate form had hampered the development of logic for centuries. Montague showed that by using type theory, that form can be preserved, if one wishes, when formalizing natural language. (In fact, he thought there were no essential differences between English and formal languages.)

Until it has been done (if it can be done), upholding (C) seems more like a pious hope.⁴

Besides, it is not obvious that the conflict would be resolved. Underspecification *reduces* ambiguity; there is no guarantee that it *removes* it. If not, the problems alluded to under (a) above might reappear: it is not enough to stipulate away the conflict. So again we might ask if there is not some other way to get, compositionally, from structure representations at some level to ordinary meanings.

(d) Claim that a *dynamic* approach to meaning avoids the conflict.

Here I am thinking above all of the work of Fernando [3], [4], which deals with the process of interpreting discourse, where ambiguity is allowed (though ambiguity at one stage may be reduced or eliminated later on in the discourse). He defines several notions of synonymy, and shows (non-trivially) that the substitution version of compositionality holds for them. This is very interesting, though it is too early to say, I think, if such an approach really reconciles (C) and (A). One reason is that Fernando so far only considers composition as concatenation of sentences and does not go into sentence structure. Another is that he starts from synonymy, and I noted earlier that it is not in general obvious that such a procedure leads to a notion of meaning which is useful for explaining understanding and communication.⁵ However, in his particular case it might very well do that, so at a more developed stage this approach might in fact succeed in handling the conflict.

(e) Admit the conflict, and conclude that one of (C) and (A) has to go.

This is the attitude of Jeff Pelletier, one of the few who have dealt at length with the issue of the relation between (C) and (A), in particular in [11]. I am sympathetic to this, though I shall want to draw a different conclusion, namely, that (C) can be adjusted so as to allow for an ambiguous semantics. How that works out occupies much of the rest of this paper. But in the next section we will look at two recent quotes on the subject, one of which shows in more detail what Pelletier is claiming.

4 Two Examples

The following comes from the Introduction to a recent collection of papers on computational semantics:

⁴Or, all right, a methodological principle, but in the sense of something you would *like* your favourite theory to satisfy, though it might not.

⁵Any synonymy satisfying the substitution version of compositionality trivially gives rise to a compositional semantics with that associated synonymy, namely, the *equivalence class semantics*, where the meaning of an expression is the set of expressions synonymous with it. This semantics, however, is a set theoretic construct, and not obviously useful for practical purposes.

At this point we may also note a strange aspect of the principle of compositionality that we did not yet consider. It speaks of “*the meaning* of a compound expression”. The use of the singular *the meaning* is common in formulations of the compositionality principle, but clearly has no basis in reality: expressions in natural language hardly ever have one single meaning. Speaking of *the meaning* is reasonable only when applied to *utterances*, where often only one of the many possible meanings of the sentence is contextually possible or relevant. This is why people can use language without constantly dealing with millions of possible meanings. (But as already noted, the meanings of utterances by their very nature do not obey Compositionality.) (Bunt and Muskens [1], pp. 15–16.)

I admit that it is not completely clear to me what the authors are saying here. It *seems* that they are saying that to apply (C) we must disambiguate, which leads us to consider utterances instead of sentences, but that for utterances (C) is false.⁶ That would amount to a *reductio* of (C). But elsewhere in the Introduction (C) is thought to be a very important principle, albeit a methodological one.

In any case, the quote relates to points (a) and (c), and possibly (e), above, as well as the remarks in section 2.2. If nothing else, I hope it illustrates the fact that the relationship between (C) and (A) is in some need of clarification.

The second quote is from Pelletier. He first considers lexical ambiguity, say in an example like

- (1) Linda approached the bank.⁷

And this type of (sentential) ambiguity is not seen as jeopardizing compositionality, for it is still felt that the meaning of these kinds of sentences (where the meaning is now interpreted as a *set* of unambiguous meanings) is a function (only) of the meaning of its parts and their manner of syntactic combination. The basic, atomic parts are allowed to have more than one meaning, and this permission is then passed up to more complex phrases containing such ambiguous parts. (Pelletier [11])

Several interesting things are being claimed here, or rather claimed, I think correctly, to be widely held:

- Lexical ambiguity is *harmless*: it poses no serious threat to (C). Call this the ‘harmlessness claim’. It could mean that there is some version of (C) adapted to lexical ambiguity, or that such ambiguity can be eliminated or ignored in connection with (C).

⁶Since the parts of an utterance are not sufficient to determine its meaning: a lot of context is needed as well. In my opinion this is not self-evident either, but I won’t pursue it here.

⁷The examples (1) – (3) are not from Pelletier [11], but similar to the ones he uses.

- Set semantics is a way to maintain (C). We already noted that this is a common idea.
- The ‘passing-up claim’: The various meanings of lexically ambiguous items are passed up to complex expressions.

These claims are of course interrelated, but it is useful to distinguish them as I have done here. Next, Pelletier discusses structural ambiguity. He argues, rather convincingly in my opinion, that while there are clear cases of structural ambiguity, such as

(2) He saw her duck under the table.

where there obviously are two different structures with the same surface form,⁸ there are also clear cases of sentences with just one structure which are still (non-lexically) ambiguous. An example would be

(3) Most critics reviewed two films.

And for these,

what compositionality *cannot* admit is that there be *no* lexical ambiguity, there be but *one* syntactic structure, and yet there be two (or more) meanings for that item. ([11], his italics.)

For lack of a better word, let me call such examples cases of *essential* ambiguity. Though Pelletier does not discuss them, one might consider underspecified meaning representations to be of this kind too: they have one structure but several meanings. Of course, we know how to represent each of these meanings, as we do for the two readings of (3), but independent arguments might show that no further syntactic structure is relevant for explaining understanding or communication.

5 Relational Semantics

I already said that an obvious move at this point is to regard a semantics for E as a *relation* R between expressions and meanings, i.e., as a subset of $E \times M$. So let us do that from now on. The non-ambiguous case is then when $R = \mu$ is a *single-valued* relation. For an expression p , let

$$R_p = \{m : R(p, m)\}$$

be the set of meanings of p . The notion of meaningfulness (belonging to the domain of R) is as before, but the notion of synonymy is no longer clear.

⁸Is there also lexical ambiguity? A very simple-minded notion of lexicon might imply that, whereas a more sophisticated one would equip lexical items with syntactic categories and perhaps lots of other features, thus making the verb *duck* and the similar-looking noun distinct atoms to begin with. Nothing I say here turns on this issue.

Let an *R-synonymy for E* be any reflexive, symmetric, and transitive relation on the domain of *R*. Now, when are two ambiguous expressions synonymous? One answer is the following:

$$p \equiv_R q \text{ iff } R_p = R_q$$

This is a natural strong notion of synonymy for ambiguous expressions.⁹ A *much* weaker notion is:

$$\begin{aligned} \equiv_{R^f} = & \text{ the transitive closure of the relation that holds between } p \text{ and } q \\ \text{iff } & R_p \cap R_q \neq \emptyset \end{aligned}$$

So we have $p \equiv_R q$ iff p and q have exactly the same meanings, whereas $p \equiv_{R^f} q$ iff there are terms p_0, \dots, p_k ($k \geq 1$) such that $p = p_0$, $p_k = q$, and $R_{p_i} \cap R_{p_{i+1}} \neq \emptyset$ for $i < k$.¹⁰ And there may other possible synonymies ‘in between’. Note that both \equiv_R and \equiv_{R^f} reduce to \equiv_μ when $R = \mu$ is single-valued.

Now, the set semantics corresponding to R , i.e., the semantics μ_R given by

$$\mu_R(p) = R_p$$

is single-valued, and its associated synonymy is \equiv_R . Hence, $\text{Rule}(\mu_R)$ holds iff $\text{Subst}(\equiv_R)$ holds (given the Domain Rule). But that’s not what concerns us at the moment. What we’d like to know is if there is some version of (C) for R .

6 A Proposal

6.1 The Idea

Consider again (3). We may assume its structure can be represented thus:

$$(4) \ \alpha(\beta(\text{most}, \text{critic}), \gamma(\text{review}, \beta(\text{two}, \text{film})))$$

We can also assume, for simplicity, that all the proper constituents of this expression are unambiguous. So the ambiguity ‘arises’ with the application of the last syntactic rule, α .

But then, isn’t it sort of obvious that in this case there is one syntactic rule but *two* corresponding semantic operations, i.e., two ways to get from the

⁹Stronger notions are conceivable, for example, variants of Carnap’s *intensional isomorphism*, where p and q have not only the same meanings but also the same structure, with constituents which also have the same meanings, etc. See Pagin [9] for a treatment of compositionality issues in such contexts.

¹⁰‘ f ’ here is meant as a mnemonic for ‘family resemblance’ in a Wittgensteinian sense: p family-resembles q if there is a family of resemblances to which both belong. This does not necessarily mean that p and q have something in common, but p has something in common with p_0 , which has something in common (not necessarily the same thing) with p_1, \dots , which has something in common with q . This gives indeed a very weak notion of synonymy; we may, for example, easily have $p \equiv_{R^f} q$ even though $R_p \cap R_q = \emptyset$, i.e., even though p and q have no meanings in common.

meanings of the parts to the meanings of the whole? To figure out the meaning of the sentence (say, as used on a particular occasion), one has to have the meanings of the parts, and furthermore to know or be able to follow these two operations, and finally to ‘choose’ between them. This still seems pretty compositional to me. Each meaning of the sentence is just as calculable from the meanings of the parts as in the non-ambiguous case. The fact that there are different ways of doing this calculation does not destroy compositionality, as long as these ways are *specified in advance and depend, as before, only on the relevant syntactic rule*.

We may then say that the meanings of the complex expression are still *determined* by the meanings of the parts and the mode of composition. The choice between the two meaning operations can be seen as a distinct *source* of ambiguity, on a par with, say, lexical ambiguity. There is no principled difference, it seems to me, between choosing an appropriate meaning for *bank* in order to understand (1), and choosing an appropriate meaning operation corresponding to α in order to understand (3). We can allow both sorts of ambiguity and still retain the basic idea behind compositionality.

Thus, while Pelletier claims that only lexical ambiguity is compatible with compositionality, I want to claim that essential ambiguity is compatible with it too. Of course I am required to formulate a version of (C) that substantiates this. But, strictly speaking, so is Pelletier for the lexical case. What we really need is a version of (C) that allows *both* kinds of ambiguity, *and* explains the particularity, and possibly the ‘harmlessness’, of the lexical variant. So that is what I will try to do.

6.2 Details

The intuitive idea behind generalizing the rule version of (C) to a relational semantics has just been roughly presented, but actually there are some niceties involved in finding the adequate precise version. The original (C) is about the existence of functions, one for each syntactic rule, and we have already seen that some take this to be a virtually empty demand. I disagreed about that, but now we are to allow several functions for each rule. Then the threat of triviality is certainly greater. The precise version of the following, for example, can be shown to be empty, in the sense of being true under practically all imaginable circumstances:

- (*) For each syntactic rule α there is a finite number of semantic operations $r_\alpha^1, \dots, r_\alpha^k$ such that any meaning of $\alpha(p_1, \dots, p_n)$ results from applying some r_α^j to some of the meanings of p_1, \dots, p_n .

Here is a precise statement of the emptiness of this notion. Call a semantics R *bounded* if for each syntactic rule α there is a natural number k such that for all $(m_1, \dots, m_n) \in M^n$, the number of meanings of grammatical terms of the form $\alpha(q_1, \dots, q_n)$ with $R(q_i, m_i)$ for $1 \leq i \leq n$ is bounded by k . That is,

$$|\bigcup\{R_{\alpha(q_1, \dots, q_n)} : \bigwedge_{1 \leq i \leq n} R(q_i, m_i)\}| \leq k$$

Then one may prove the following:¹¹

Fact 1

If R is bounded (and the Domain Rule holds) then () is satisfied.*

There seems to be no argument why any reasonable semantics should not be bounded, and so in this sense, (*) is a much too weak condition.

I should mention at this point that the previous claim depends on one assumption that I am making throughout here:

(fin) Expressions have a finite number of meanings, i.e., each R_p is finite.

If some complex expression has an infinite number of meanings but its constituents only a finite number of meanings, then (*), and the conditions discussed below, are simply false. This means that relational semantics is in reality intended for notions where the the number of meanings is relatively small. So it is not intended for referential ambiguity, if meaning is taken to be reference, so that a pronoun might then have infinitely many meanings, or a quantified sentence might have have infinitely many meanings depending on the universe of discourse. Of course, a semantics might handle such referential indeterminacy without construing it as ambiguity of meaning.

In view of the problem with (*), one might try to strengthen the requirement by changing an implication to an equivalence:

(**) For each syntactic rule α there is a finite number of semantic operations $r_\alpha^1, \dots, r_\alpha^k$ such that m is a meaning of $\alpha(p_1, \dots, p_n)$ if and only if m results from applying some r_α^j to some of the meanings of p_1, \dots, p_n .

But then it becomes too strong: One can find reasonable semantics for which (**) fails.¹² So I will not discuss these last conditions further, but instead go directly to the version of (C) for a relational semantics R that I'd like to propose.

The remedy, however, is to have one version of (C) for each k ; the natural number k measures the 'degree of essential ambiguity' allowed:

Rule^k(R) For each syntactic rule α there are semantic operations $r_\alpha^1, \dots, r_\alpha^k$ such that for each $m \in R_{\alpha(p_1, \dots, p_n)}$ there is some j and there are $m_i \in R_{p_i}$, $1 \leq i \leq n$, such that

¹¹This and other results in this paper can all be proved rigorously in a precise version of the present framework, as in Hodges [8] and Westerståhl [14]. I save the details of these proofs for another occasion.

¹²Suppose m is a meaning of $\alpha(p_1, \dots, p_n)$, with

$$m = r_\alpha^j(m_1, \dots, m_n),$$

where $m_i \in R_{p_i}$. Suppose further that m_1, \dots, m_n also are meanings of expressions q_1, \dots, q_n , respectively, and that $\alpha(q_1, \dots, q_n)$ is meaningful. Then it would follow from (**) that m has to be one of the meanings of $\alpha(q_1, \dots, q_n)$. But that cannot hold in general: though the p_i and q_i share *one* meaning, there seems to be no reason that $\alpha(p_1, \dots, p_n)$ and $\alpha(q_1, \dots, q_n)$ share the *particular* meaning m .

- (a) $m = r_\alpha^j(m_1, \dots, m_n)$,
- (b) for each j' , $1 \leq j' \leq k$, $r_\alpha^{j'}(m_1, \dots, m_n) \in R_{\alpha(p_1, \dots, p_n)}$.

Two quick comments. First, note that when $k = 1$ there just one semantic operation per rule, as in the standard version of (C), but there may still be ambiguities of the non-essential kind, e.g., lexical ambiguities. We'll come back to this in section 7.

Second, you may wonder where condition (b) comes from. I submit that my main reason for including it is that it makes the following hold:

Fact 2

When $R = \mu$ is single-valued, we have, for each $k \geq 1$, that $\text{Rule}^k(\mu)$ is equivalent to $\text{Rule}(\mu)$.

For, if there is just one meaning m in $R_{\alpha(p_1, \dots, p_n)}$, and just one m_i in each R_{p_i} , then (b) forces it to be the case that $m = r_\alpha^1(m_1, \dots, m_n) = \dots = r_\alpha^k(m_1, \dots, m_n)$, so $r_\alpha^1 = \dots = r_\alpha^k$. But without (b) this may fail. And I think it is a good criterion for any reasonable version of (C) for ambiguous semantics that it reduces to standard compositionality in the non-ambiguous case.

Next, I want to say a few words about some further properties of the proposed form of (C) for relational semantics.

6.3 Substitution Versions?

$\text{Rule}^k(R)$ generalizes the rule version of compositionality. Is there an equivalent substitution version also for ambiguous semantics? I think not. Note first that $\text{Rule}^k(R)$ clearly implies $\text{Rule}^{k+1}(R)$ for all k . Then, we observe the following

Fact 3

$\text{Rule}^1(R)$ does not imply $\text{Subst}(\equiv)$, for any synonymy \equiv for E between \equiv_R and \equiv_{R^f} .

This can be seen by the following abstract example. Suppose

$$\begin{aligned} R_a &= R_b = \{m_1, m_2\} \\ R_{\alpha(a)} &= \{m\} \\ R_{\alpha(b)} &= \{m'\}, \end{aligned}$$

where all these meanings are distinct. Then $\text{Rule}^1(R)$ can hold, with $r_\alpha^1(m_1) = m$ and $r_\alpha^1(m_2) = m'$. On the other hand, $a \equiv_R b$ by assumption, but $\alpha(a) \not\equiv_{R^f} \alpha(b)$: there is not even a ‘family resemblance’ (cf. note 10), we are assuming, between m and m' . So $\text{Subst}(\equiv)$ must fail for any synonymy \equiv between \equiv_R and \equiv_{R^f} .

But couldn't there be some ‘substitution-like’ version even if it doesn't have exactly the form $\text{Subst}(\equiv)$? I think the following observation makes that unlikely. Consider a condition $\text{Rule}^2(\mu)^-$, for an unambiguous semantics μ , which is just like $\text{Rule}^2(\mu)$ as defined above, *minus* the condition (b). That is, for each

syntactic rule α there are two operations r_α^1 and r_α^2 such that whenever $\alpha(p)$ is meaningful (let us assume that α takes just one argument, for simplicity), $\mu(\alpha(p))$ is equal to *either* $r_\alpha^1(\mu(p))$ *or* $r_\alpha^2(\mu(p))$. Never mind that this seems like a very *ad hoc* condition; the question is if we can find a substitution version. Here it is:

Fact 4

Rule²(μ)⁻ is equivalent to the following condition (in the case of a 1-place α): If $p \equiv_\mu q \equiv_\mu r$, and $\alpha(p), \alpha(q), \alpha(r)$ are all μ -meaningful, then at least two of them are μ -synonymous.

It is clear that if all of $\alpha(p), \alpha(q), \alpha(r)$ have distinct meanings, then these three meanings cannot be obtained by two functions applied to the same argument (i.e., the argument $m = \mu(p) = \mu(q) = \mu(r)$). With a little work one can see that the implication goes in the other direction as well.

My point here is just to illustrate how complicated it becomes to find substitution versions when more than one meaning operation correspond to each syntactic rule.¹³ Our condition Rule²(R) included clause (b), but that doesn't seem to make things any simpler. In sum, it would appear that there is no sensible substitution version of Rule^k(R).

But perhaps there is no problem with that conclusion. Tentatively, I would interpret it in the following way: The intuitions behind the rule version and the substitution version of (C) are really quite different. The two versions are extensionally the same for non-ambiguous semantics, but not intensionally. It is more like an accident that they coincide in that case. And while there is a quite straightforward way to generalize the rule version to a relational semantics, there is no corresponding generalization of the substitution version. In the rule case we can still talk about semantic operations corresponding to syntactic rules. In the substitution case we would, to begin with, need an adequate notion of synonymy. But maybe there just is no natural candidate for such a notion when expressions are allowed to have more than one meaning.¹⁴

6.4 Non-triviality

In contrast with the condition (*) mentioned in section 6.2, Rule^k(R) is a substantial requirement in the sense that it is easy to falsify. That is, it is easy to define semantics for which Rule^k(R) fails. In fact, it suffices to make some assumptions concerning the *number* of meanings of certain expressions. Roughly, if there are too many meanings of a complex expression compared to the number of meanings of its parts, then Rule^k(R) cannot hold. This should be rather

¹³Note also that the condition in Fact 4 deals only with immediate constituents, not arbitrary ones as in the substitution versions mentioned earlier. A formulation in terms of arbitrary constituents would be even more complicated.

¹⁴More precisely, the claim is that *given* a notion of meaning (a semantics), there might be no interesting *corresponding* notion of synonymy. As I mentioned (section 3), Fernando [3], [4] proposes to start directly with synonymy relations and dispense, in a sense, with meanings. That is a different approach from the one pursued here, and those synonymies might well be both natural and interesting.

obvious; I include the following precise statement just for the record. Here, if X is a set, $|X|$ is its cardinality.

$\text{Card}^k(R)$ For any (n -ary) syntactic rule α , and any selection of n -tuples of expressions p_{i_1}, \dots, p_{i_n} ($i \in I$) that α can be applied to,

$$|\bigcup_{i \in I} R_{\alpha(p_{i_1}, \dots, p_{i_n})}| \leq k \cdot |\bigcup_{i \in I} (R_{p_{i_1}} \times \dots \times R_{p_{i_n}})|.$$

As a special case, if $\alpha(p_1, \dots, p_n)$ has more meanings than k times the number of n -tuples of meanings in $R_{p_1} \times \dots \times R_{p_n}$, then $\text{Card}^k(R)$ fails, and, obviously, so does $\text{Rule}^k(R)$. It is not hard to show the following:¹⁵

Fact 5

$\text{Rule}^k(R)$ implies $\text{Card}^k(R)$.

7 On Lexical Ambiguity

7.1 ‘Passing-up’ and Set Semantics

Now let us see how the present proposal handles lexical ambiguity. Recall the claims made in the quote from Pelletier in section 4 above: the ‘harmlessness claim’, the ‘passing-up claim’, and the claim that set semantics will do the job.

The last claim can be laid to rest in the following way. It is not a new point; it is made also in van Deemter [12] (p. 208), with an example similar to the one used here (which is a variant of the example used in section 6.3). Suppose

$$\begin{aligned} R_a &= R_b = \{m_1, m_2\} \\ R_{\alpha(a,a)} &= \{m, m'\} \\ R_{\alpha(b,b)} &= \{m\}, \end{aligned}$$

The idea is that while a and b have the same two meanings, repetition of a (in context α) ‘passes up’ both of these (with the natural constraint that both occurrences of a have to mean the same, so that there are 2, not 4, meanings of $\alpha(a, a)$), whereas repetition of b only ‘passes up’ one. There seems to be nothing strange, in principle, about this kind of example.¹⁶ Then $a \equiv_R b$, but $\alpha(a, a) \not\equiv_R \alpha(b, b)$, so $\text{Subst}(\equiv_R)$ fails. But this means, as already noted, that $\text{Rule}(\mu_R)$ fails too.¹⁷

So the set semantics in this case is not compositional. In other cases it might be. For example, ‘Cooper storage’ (see Cooper [2]), which is a version of set

¹⁵Jouko Väänänen pointed out to me by means of an example that the converse implication is far from true.

¹⁶Or so it seems to me. I have realized, however, that many people feel the example strongly contradicts basic intuitions concerning compositionality. The matter merits further discussion, but I will not pursue it here.

¹⁷In the (very) special case when complex expressions as well as their parts are *sentences*, this is related to the failure of the so-called *disjunction semantics* to be compositional (or to do justice to our intuitions) — this is also pointed out in van Deemter [12].

semantics, is compositional. My point is just that there is no guarantee that switching to a set semantics will rescue compositionality.

But even if the set semantics above is not compositional, $\text{Rule}^1(R)$ may well hold: there can be one operation r_α such that each meaning of $\alpha(a, a)$ and $\alpha(b, b)$ is calculable by means of it from the meanings of the constituents; say,

$$m = r_\alpha(m_1, m_1),$$

$$m' = r_\alpha(m_2, m_2).$$

The conclusion, then, is that set semantics should not be expected to be compositional. But there can still be ‘passing-up’ of meanings of parts to meanings of complex expressions. Now, if there is any ‘passing-up’ going on at all, it would have to be by means of meaning operations of the kind used in the rule version of (C). And we have seen how these can exist in the presence of ambiguity. Thus, we can stick to the ‘passing-up’ claim even if set semantics goes out the window.

Note also that there may be constraints on what is allowed to be ‘passed up’. We hinted at one example of such a constraint above. Another one is illustrated by the sentence (1), repeated below, compared with (5):

(1) Linda approached the bank.

(5) Linda robbed the bank.

In (5), only one of the meanings of *bank* is ‘passed up’. This sort of constraint has to be built into the corresponding meaning operations.

7.2 The Sense in which Lexical Ambiguity is Harmless

One idea is that if there is only lexical ambiguity, there is no problem about compositionality. But what does it mean that there is *only* lexical ambiguity? Certainly not that only lexical items are ambiguous; so are presumably certain complex expression in which they occur. A necessary condition seems to be:

(lex) If a complex expression has more than one meaning, so does at least one of its immediate constituents (and hence at least one of its lexical constituents).

However, (lex) is consistent with having, say, $|R_a| = 2$ and $|R_{\alpha(a)}| = 3$, which goes against the ‘passing-up’ idea in $\text{Rule}^k(R)$ when $k \geq 2$. Thus, what we need to do is let $k = 1$:

Fact 6

$\text{Rule}^1(R)$ implies (lex).

Therefore, I suggest that the notion of the compositionality of a semantics with only lexical ambiguity is captured adequately by the condition $\text{Rule}^1(R)$.

However, the idea that lexical ambiguity is harmless doesn't apply — or shouldn't apply — just to the case where there is only lexical ambiguity. It ought to mean, I think, that such ambiguity can always be *eliminated*, without disturbing compositionality. Of course this presupposes a notion of compositionality compatible with ambiguity, which is precisely what I have tried to present here. Let us see, finally, how the 'eliminability claim' works out.

The trick is quite familiar: Replace ambiguous atoms a with n meanings by indexed atoms a_1, \dots, a_n , each with just one of the meanings of a . Assuming the new atoms and the old one have the same surface form, this easily yields a new set of structured expressions, each of which is a *lexical disambiguation* of an old expression. Now suppose R is a semantics for the old expressions such that $\text{Rule}^k(R)$ holds. Then a new semantics R^d for indexed expressions is defined in the obvious way for indexed atoms, and inductively for complex expressions $\alpha(p_1, \dots, p_n)$ as follows:

(dis) $R^d(\alpha(p_1, \dots, p_n), m)$ iff $R(\alpha(p_1^-, \dots, p_n^-), m)$ and

$$m = r_\alpha^j(m_1, \dots, m_n)$$

for some $j \leq k$ and some $m_i \in R_{p_i}^d$, $1 \leq i \leq n$ (where p^- is the result of deleting all indices on atoms in p).

Fact 7

- (a) If $R^d(p, m)$ then $R(p^-, m)$.
- (b) If $R(p^-, m)$ then $R^d(q, m)$, for some lexical disambiguation q of p^- .
- (c) R^d has no lexical ambiguities, and $\text{Rule}^k(R^d)$ holds. In particular, if $k = 1$, R^d is a single-valued and compositional semantics.¹⁸

This, I would claim, is the sense in which lexical ambiguity is harmless, in the context of (C).¹⁹

As an example, *approach the bank* becomes *approach the bank₁* and *approach the bank₂*, with distinct and unique meanings. What about *rob the bank*? With the simple-minded grammatical changes suggested above, *rob the bank₁* and *rob the bank₂* both become wellformed, but if *bank₂* means 'river bank', the latter is not meaningful, since the appropriate meaning operation is not defined for river banks. So only the other meaning of *bank* is 'passed up'.

Consider also our previous abstract example with

$$\begin{aligned} R_a &= R_b = \{m_1, m_2\} \\ R_{\alpha(a,a)} &= \{m, m'\} \\ R_{\alpha(b,b)} &= \{m\}. \end{aligned}$$

¹⁸All of this can be stated and proved formally in (a slight extension of) the framework of Hodges [8].

¹⁹Of course, I am by no means implying that lexical ambiguity in itself is trivial or uninteresting. I am only discussing to what extent it may be a problem for compositionality.

Here we get indexed atoms a_1, a_2, b_1, b_2 with

$$\begin{aligned} R_{a_1}^d &= R_{b_1}^d = \{m_1\} \\ R_{a_2}^d &= R_{b_2}^d = \{m_2\}. \end{aligned}$$

$\alpha(a_1, a_1)$ and $\alpha(b_1, b_1)$ each means m , $\alpha(a_2, a_2)$ means m' , but $\alpha(b_2, b_2)$ has no meaning: by (dis), it cannot mean m since $r_\alpha(m_2, m_2) \neq m$, and it cannot mean m' since $\alpha(b, b)$ does not mean m' .

What about $\alpha(a_1, b_1)$? You might suspect a breach of compositionality here, in the substitution version:

$$a_1 \equiv_{R^d} b_1,$$

but

$$\alpha(a_1, a_1) \not\equiv_{R^d} \alpha(a_1, b_1),$$

since the expression on the right hand side is not meaningful. But there is not: compositionality in the substitution version only says something about the case when *both* complex terms are meaningful (cf. the condition $\text{Subst}(\mu)$ in section 2.3.2).²⁰

8 Conclusions

I have argued for the following claims:

- Compositionality — condition (C) — does not express a semantic theory. It is a general requirement on such theories, motivated by a certain view of how linguistic understanding and communications works.
- We *can* make sense of (C) in the presence of ambiguity, in particular lexical and what I called essential ambiguity. Moreover, this can be done in a way that preserves the original intuition behind the rule version of compositionality, in particular as a (necessary) condition on explanations of linguistic understanding and communication.
- The generalization — $\text{Rule}^k(R)$ — makes a distinction as to the ‘degree of essential ambiguity’. When this degree is 1, there is only lexical ambiguity. For any degree k , if R is single-valued, $\text{Rule}^k(R)$ reduces to ordinary compositionality.
- Set semantics is not the way to reconcile (C) and (A); in fact I would venture the claim that there is no reason to expect set semantics to be compositional.

²⁰However, we do have a breach of what Hodges calls the *Husserl property*, which says that synonymous terms can be substituted for each other with preserved meaningfulness. But then, the Husserl property may not hold for real languages anyway.

- With a given relational (ambiguous) semantics there might be no obvious natural notion of synonymy, hence no substitution version of compositionality.²¹ But, I suggest, this need not be a great loss, since it is the rule version of (C) that drives our idea of compositionality as a constraint on explanations of linguistic communication and understanding.
- Lexical ambiguity *is* harmless, in the sense that it is (a) compatible with the extended version of (C), and (b) eliminable while preserving this extended sense of compositionality.

An even shorter summary: I propose a generalization of the notion of compositionality which applies also to ambiguous semantics, but is the same for non-ambiguous semantics. It is not a proposal for a particular type of semantics; rather, as are all abstract versions of (C), it is a constraint on semantics in general. One hope is that the generalized notion would make the relation between compositionality and ambiguity clearer.

References

- [1] H. Bunt and R. Muskens. Computational semantics. In H. Bunt and R. Muskens, editors, *Computing Meaning, Volume 1*, pages 1–32. Kluwer Academic Publishers, Dordrecht, 1999.
- [2] Robin Cooper. *Quantification and Syntactic Structure*. D. Reidel, Dordrecht, 1983.
- [3] T. Fernando. Ambiguity under changing contexts. *Linguistics and Philosophy*, 20, 1997. 575–606.
- [4] T. Fernando. Ambiguous discourse in a compositional context. *Journal of Logic, Language and Information*, 10:63–86, 2001.
- [5] J. Groenendijk and M. Stokhof. Dynamic predicate logic. *Linguistics and Philosophy*, 14:39–100, 1991.
- [6] J. Higginbotham. Linguistic theory and Davidson’s program in semantics. In E. LePore, editor, *The Philosophy of Donald Davidson: Perspectives on Truth and Interpretation*. Blackwell’s, Oxford, 1986.
- [7] W. Hodges. Compositionality is not the problem. *Logic and Logical Philosophy*, 6:7–33, 1998.
- [8] W. Hodges. Formal features of compositionality. *Journal of Logic, Language and Information*, 10:7–28, 2001.
- [9] Peter Pagin. Communication and strong compositionality. *Journal of Philosophical Logic*, 32:287–322, 2003.

²¹But note the *caveat* mentioned in footnote 14.

- [10] F. J. Pelletier. The principle of semantic compositionality. *Topoi*, 13:11–24, 1994.
- [11] F. J. Pelletier. Semantic compositionality: Free algebras and the argument from ambiguity. In M. Faller, S. Kaufmann, and M. Pauly, editors, *Proceedings of Seventh CSLI Workshop on Logic, Language and Computation*. CSLI Publications, Stanford, 1999.
- [12] K. van Deemter. Towards a logic of ambiguous expressions. In K. van Deemter and S. Peters, editors, *Semantic Ambiguity and Underspecification*, pages 203–237. CSLI Publications, Stanford, 1996.
- [13] D. Westerståhl. On mathematical proofs of the vacuity of compositionality. *Linguistics and Philosophy*, 21:635–643, 1998.
- [14] Dag Westerståhl. On the compositional extension problem. *Journal of Philosophical Logic*, 33:549–582, 2004.
- [15] W. Zadrozny. From compositional to systematic semantics. *Linguistics and Philosophy*, 17:329–342, 1994.