

## Remarks on scope ambiguity\*

*Dag Westerståhl*

*Department of Philosophy, Göteborg University*

### **Abstract**

*It has been argued that some scope ambiguities are not really structural: the underlying phrase structure is ambiguous, and no further structure can be independently motivated. Assuming this, with a simple case of quantifier scoping as illustration, we look at the consequences for semantics, and in particular for the principle of compositionality. I observe that attempts to save the usual version of compositionality by revising the notion of meaning (in particular by taking the new meaning to be the set of old meanings) are not guaranteed to work. My main claim is that if one accepts instead a relational semantics, where structured as well as lexical*

---

\* Jens Allwood is a linguist with very broad philosophical interests. We have been colleagues for an inordinate number of years, and still meet regularly to talk about linguistic and philosophical matters. In the context of the present paper, I recall that Jens was one of the first to introduce Montague Grammar in Sweden, in the 70's. I think he would agree with my claim that Montague's treatment of scope ambiguity is not wholly convincing. Since those early days, however, our approaches to our common interest – language – have increasingly diverged, and I am afraid he might not agree with much else I say here. But then, Jens rarely agrees with much of what anyone says.

*expressions are allowed to have more than one meaning, the compositionality principle can be adapted without loss of explanatory power.*

**Keywords:** scope, ambiguity, compositionality, relational semantics

## **1. Introduction**

Since Montague we have grown accustomed to the idea that structural ambiguities, and in particular ambiguities of scope, always have a syntactic counterpart. To every syntactic rule  $\alpha$  corresponds a unique semantic rule  $r_\alpha$ , and since there is nothing more to semantics than rules of this form and the input from the lexicon, all meaning variation in a surface string goes back to either lexical ambiguity or differences in the syntactic derivation of the string. Lexical ambiguity can be dealt with by simply introducing different terms for different meanings of a lexical item, and so all ambiguities are nicely disposed of.

This is an elegant picture, and it was no trivial matter at the time to see that it could be rigorously carried through for large parts of natural languages. Today we should perhaps be less impressed by this particular aspect of Montague's achievement; it has become almost obvious that if one *wants* to account for ambiguity in this way, it is *possible* to do it. So what are the *reasons* for choosing this or that syntactic format? Specifically, if the *only* reason for introducing certain syntactic rules is to account for semantic ambiguity, perhaps that reason is not good enough?

In this short and programmatic paper I shall not discuss syntax, but only indicate, by means of a simple example, some consequences of abandoning Montague's idea and allowing syntax to 'underdetermine' semantics, even as regards some issues of scope. Such a move is rather natural and has been suggested many times in the literature. But what has not been discussed with equal care, I think, is the methodological consequences that ensue. In particular, how does this move affect another basic principle of semantics, that of *compositionality*?

## 2. The example

The example is a basic case of quantifier scope ambiguity, but rather than the philosopher's standard "Every man loves a woman" I choose a slightly more interesting (and realistic) version. Consider the following two sentences:

- (1) Two countries objected to every proposal.
- (2) Two knives belonged to every equipment.

I will make a number of assumptions concerning (1) and (2). The first two should be completely uncontroversial.

(a) To understand what these sentences *say* (e.g. to decide whether they are true or not) one needs a *context*: an utterance situation, some worldly background, etc. However, 'default' contexts easily spring to mind, allowing one to rather effortlessly make sense of the sentences. For example, (1) could be about an assembly of representatives of nations, such

as UN or EU, discussing measures to adopt on some sensitive matter. And (2) could be about the equipments issued to a group of boy scouts before one of their hikes. (Think of (2) uttered after the hike by the leader of the (say) 10 boys, who goes on to say: “But so far only 19 knives have been returned!”)

(b) In these ‘default’ contexts, the determiners *two* and *every* have different scope in (1) and (2). (1) (normally) says that there were two particular countries whose representatives opposed every proposal made during the assembly, *not* that for every such proposal there were two countries (possibly different ones for different proposals) opposing it. But (2) (normally) does *not* say that there were two specific knives that somehow belonged to each of the equipments. Rather, it (normally) says that each equipment contained two knives (most likely different knives for different equipments).<sup>1</sup>

I think (a) and (b) are perfectly obvious. Nevertheless, it is sometimes argued that quantifier scope ambiguities are rare and that the surface syntax provides a default reading. (1) and (2) should make one suspicious of such claims.<sup>2</sup> But that suspicion needs to assume something about the surface syntax. I think (almost) no one would object to the following:

(c) (1) and (2) have the *same* surface syntax.


---

<sup>1</sup> Furthermore, it is easy to imagine a context where (1) has the other, ‘inverse scope’ reading. That reading is the default for (2) but, (with greater effort) one can imagine a situation where the other reading was intended instead.

<sup>2</sup> Note that the claim does *not* follow in any straightforward way from the (in itself interesting) fact that speakers normally and without thought pick out one of the scopes in actual utterances. The issue is not the ambiguity of utterances but of the *sentences* they use.

The next assumption might look more debatable, but it really isn't, since it is meant to be independent of any particular format for syntax. Pick your favorite syntactic framework: the odds are overwhelming that (a suitably adapted version of) (d) below will hold:

(d) The basic surface structure of both (1) and (2) can be represented as follows.



The labels here have no importance; nor does it matter which phrases are taken to be ‘heads’ and which are ‘specifiers’, nor whether additional (X-bar style) labels are inserted. We are only talking about basic *structure*.

Now, we all know how to endow T with additional structure in ways that give (1) and (2) *different* syntactic analyses, corresponding to the different scopes of the determiners. For example, we may allow transformations (movement, quantifier raising) to turn T into one of two more complex trees (logical forms) where scope is fixed, and argue that meaning is only assigned to these latter forms. Or we may have syntactic rules that generate the more complex forms directly, without using T, or perhaps using T for

one of the two scopes (the default scope of (1)), and a more complex form for the other.

My last assumption, (e) below, is the only controversial one in the present context. It seems to me (and it has seemed to others, for example, Pelletier 1999) that the *only* reason to introduce such additional structure is to account for the ambiguity. In terms of syntax alone, T is all the structure there is. That is, in this case there is no syntactic structure beyond phrase structure:

(e) The syntactic structure of (1) and (2) is (essentially) T.

I think one can make a pretty good case for (e), but this will not be pursued here. The point is rather the following: *Assuming* (a) – (e), what are the consequences?

### **3. Consequences for compositionality**

The principle of compositionality says that the meaning of a complex expression is determined by the meanings of its parts and the ‘mode of composition’ (usually, the syntactic rule used to derive the expression from its immediate parts). This presupposes that each expression has (at most) one meaning, or, as Montague put it, that expressions have been *disambiguated*. Two ways of obtaining such disambiguation were sketched above, and in both cases classical compositionality can be maintained.<sup>3</sup>

---

<sup>3</sup> Barker and Jacobson (2007) call the first option (which is consonant with generative grammar, for example in many of the forms it has been given by Chomsky during the

However, if our example above is to be taken seriously, it may seem that we simply have to give up compositionality. Pelletier (1999) draws this conclusion, taking the occurrence of non-lexical ambiguities not explainable by differences in syntactic derivation as evidence that the principle of compositionality fails. But there are at least two other options available.

The first is to revise the notion of meaning. Standardly, the meaning of a sentence is its truth conditions, often taken to be embodied in a *proposition* recording its truth value under various circumstances or possible worlds. Analogously, the meanings of parts of sentences are entities that suitably contribute to propositions. One idea, then, is to take meanings to be something *less specific*, so that when the semantic analysis is done, the meaning of the whole sentence is not yet a proposition but something more abstract, which *yields* a proposition by some further operation, perhaps a pragmatic one. Applying the same idea to parts of sentences, each syntactic unit will still have a unique meaning. Of course, such a maneuver does not guarantee compositionality, but it removes one obstacle to it.

In fact, this move is quite familiar, although the reasons for making it have usually not been tied to a wish to uphold compositionality (indeed, the move has sometimes been tied to a rejection of compositionality). Instead,

---

years) *indirect* compositionality, since transformations are syntactic rules without semantic counterparts. The second option on the other hand (which is consonant with Montague Grammar), is *direct*. This issue does not really concern compositionality but rather what computer scientists have called the *Domain Principle*, i.e. that sub-expressions of meaningful expressions are also meaningful. Hodges (2001) notes that compositionality has natural formulations in either case (equivalent under DP; see section 8 below). Dowty (2007) points out that the distinction is not absolute, since an indirect semantics can often be recast as a direct one.

the motivation has been, roughly, that semantics should not just give us the means to derive all possible propositions expressed by a sentence, but also say something about how speakers usually arrive at *one* of these propositions. For Montague, this belonged to *pragmatics*, but the semantics-pragmatics borderline is something that has been much debated recently. However the line is drawn, a marked tendency has been that semantics should become more relevant to actual *understanding and communication* than Montague's framework allowed.

I will have something brief to say about this kind of move and its relation to compositionality in section 6, but my main interest here is in the second option referred to above. This idea is simply to give up, not compositionality, but the assumption that every expression has a unique meaning. That is, the idea, which is rarely carried out in the literature, is to let meaning association be a *relation* rather than a *function*, and then to attempt a reformulation of the compositionality principle for relational semantics.<sup>4</sup> What such an attempt might look like will be sketched in section 8.

#### **4. Why compositionality?**

At this point the reader might wonder why I am making such a fuss about compositionality. Why do I seek to uphold compositionality in the framework of a relational semantics, and why am I worried (see section 6)

---

<sup>4</sup> One of the rare exceptions is Debusmann et al. (2004), who give a relational grammar in this sense for a fragment of English (in a Dependency Grammar format), and for partly the same reasons as here. They do not discuss compositionality, however. Thanks to Joakim Nivre for bringing this work to my attention.



that those who advance the first option mentioned above are not sufficiently concerned with upholding it?

My answer here must be brief. I am simply adhering to a tradition which can be said to begin with Frege, and which is largely adhered to among linguists and philosophers of language today, according to which there are certain *remarkable facts* about human language that need to be *explained*, and compositionality is an important ingredient in the explanation. These facts are, roughly, that speakers of a language are able to effortlessly understand sentences they have never heard or come across before, sentences that moreover express thoughts (propositions) that are likewise *new* to the speaker. This is often called the *productive* aspect of language use. It should be contrasted with what can be called its *creative* aspect. The latter concerns speakers' ability to invent new linguistic constructions (fresh idioms and metaphors are examples), thus creating new semantic rules. It is remarkable too, I think, that speakers are able to do this with such apparent ease, and that others are often able to correctly 'guess' (if that is what they do) how these constructions and rules are to be understood. But productivity does not rely on creativity. On the contrary, its explanation in terms of compositionality is precisely that the *given* (syntactic and semantic) rules of language are used, and that it is the knowledge of these rules, together with lexical knowledge, which allows the speaker to immediately grasp applications of these *same* rules in ways she has not done or come across before, although they are implicit in the rules.<sup>5</sup>

---

<sup>5</sup> The number of ways in which these rules can be applied is infinite, or at least very large (even with a small lexicon), *much* larger than the instances a speaker actually

It must be added that all of the above has been denied. Philosophers and (to a lesser extent) linguists and logicians have denied that Frege expressed or believed in compositionality, or that the linguistic facts referred to exist, or that they are remarkable or in need of explanation, or that compositionality can explain them or indeed explain anything at all. But this is not the place to review these arguments. The position I am taking is not deviant, indeed it is mainstream, and here I will simply presuppose that there is a value in maintaining compositionality.<sup>6</sup>

## 5. Compositionality and context

Before coming to the point a word must be said about context. That natural language is strongly context-dependent is a familiar fact, though recently there has been much debate about how this is to be accounted for, whether it belongs to semantics or pragmatics, etc. Most of the people involved agree at least that certain basic indexicals – like *I*, *you*, *here*, *tomorrow* – are context-dependent in a very systematic way, largely determined by linguistic rules. It is thus reasonable to incorporate such context-dependence in the corresponding *meanings* (for example, by letting the new meanings be functions from contexts to the old meanings). But doing that has consequences for compositionality. Roughly, there are now additional ways that the compositionality of a proposed semantics might *fail*: not only may substitution of synonymous parts fail to preserve meaning, but also

---

comes across during her lifetime. Note, furthermore, that I am not saying that compositionality could not play a role in explaining linguistic creativity too.

<sup>6</sup> For recent overviews of this debate, see Szabó (2007) or Pagin and Westerståhl (to appear).

what Pagin and Pelletier (to appear) call *context shift* may have that effect. Put differently, there are distinct versions of the principle of compositionality, depending on whether one composes meanings *before* or *after* the contextual contributions are in place.

But I mention this only to disregard it. Here I want to focus on just one particular phenomenon – how certain kinds of ambiguity affect compositionality – and therefore I omit semantic features not immediately relevant to this issue. It is possible to bring out the effects of ambiguity without invoking context-dependence. However, I want to emphasize that in a fuller treatment, context-dependence should *not* be disregarded. In fact, context-dependence and ambiguity are *related* in interesting ways that need to be accounted for, but that must be left for another occasion.

## **6. Set semantics and other forms of underdetermined meanings**

As I said, one way to restore compositionality in the face of ambiguity is to revise the notion of meaning. Perhaps the simplest idea here is to use the *set* of possible meanings instead. Suppose meaning assignment is a *relation*  $R$  between (possibly analyzed) expressions and meanings, rather than a single-valued *function* from expressions to meanings. For each expression  $e$ , let

$$R_e = \{m: R(e,m)\}$$

be the set of meanings of  $e$ . (Assume a set  $E$  of expressions and a set  $M$  of meanings to be given.) Now introduce a new semantics  $\mu_R$  by defining, for each  $e$  in  $E$ ,

$$\mu_R(e) = R_e$$

This semantics is single-valued – each expression has a unique meaning – and so the issue of its compositionality may be raised in the standard way.

This way to handle ambiguity is quite common. For example, Pelletier (1999) indicates that lexical ambiguity can be treated by giving each lexical item the set of its various meanings as semantic value (new meaning), and these values are then ‘passed up’ to complex expressions in a natural way. Likewise, Cooper storage (see Cooper 1983) can be seen as a more sophisticated version of the same idea, this time not for lexical items but for ambiguities of scope. Furthermore, an intuitively appealing propositional interpretation of  $R_e$ , in the special case when  $e$  is a sentence, is simply as the *disjunction* of its members.

A lot can be said about this idea and its manifestations in the literature, but here I will make just one comment: In general, *there is no principled guarantee that the set semantics is compositional*, even when the original semantics is.

This may seem puzzling. Suppose a binary syntactic operation  $\alpha$  combines ambiguous lexical items  $a$  and  $b$ , and that a corresponding semantic operation  $r_\alpha$  combines their meanings. Can’t we then just define a new

semantic operation that operates on the sets of meanings of  $a$  and  $b$ , respectively, and yields the set of all  $r_\alpha(m_1, m_2)$ , where  $m_1$  is a meaning of  $a$  and  $m_2$  is a meaning of  $b$  (and  $r_\alpha(m_1, m_2)$  is defined), as output? The problem is that there may be all kinds of constraints on when the combination is meaningful. The simple-minded ‘lifting’ of  $r_\alpha$  just indicated may give a (set) meaning to the combination when it shouldn’t, and if we are unlucky, this can result in an irreparable failure of compositionality. For example, it could happen that  $\alpha(a, b)$  is meaningful but  $\alpha(a, c)$  isn’t, or that one of the meanings of  $\alpha(a, b)$  is not available for  $\alpha(a, c)$ , even though  $b$  and  $c$  have the *same* meanings. If so, our ‘lifting’, and indeed any set semantics, will give  $\alpha(a, b)$  and  $\alpha(a, c)$  *different* meanings, violating compositionality.<sup>7</sup>

This example is abstract but the point, I think, is general. Sets of ‘old’ meanings are rather different objects than the ‘old’ meanings themselves, and it is not obvious that natural operations on ‘old’ meanings straightforwardly lift to equally natural operations on sets, or indeed to any operations on sets. My point, therefore, is that although switching to set semantics (or its variants) looks like a natural way to get around ambiguity, the issue of compositionality is not thereby automatically resolved; on the contrary it needs to be approached anew.

---

<sup>7</sup> The counterexample – which is similar to one given in van Deemter (1996) – involves failure of what Hodges (2001) calls the *Husserl property*, roughly that intersubstitution of synonymous expressions preserves meaningfulness. It is debatable whether natural languages have this property. In any case, it is not hard to find natural proposals for semantics that violate it.

## 7. Relational semantics

We now come to what (to me) looks like the most straightforward way to deal with ambiguity, especially ambiguity that is neither lexical (or induced by lexical ambiguity) nor structural, as in our example from section 3.<sup>8</sup> The idea is: Accept that the semantics (the meaning assignment) is relational! *Prima facie*, the fact that some sentences, even when their syntactic structure is taken into account, have more than one meaning doesn't seem to be a great obstacle to understanding and communication. The hearer needs to pick up the reading that the speaker intended, but this could be achieved in a number of ways. It is not at all clear that imposing a more complex underlying structure is especially helpful. After all, with lexical ambiguity, as in *John went to the bank* (when both a river bank and a financial bank are 'salient' in the utterance situation), the hearer has to rely on other cues to get at the right interpretation, and there is no reason, it seems to me, that the same could not hold for quantifier scope ambiguity.

Our question, then, is whether such a relational semantics destroys the idea of compositionality, or the notion that compositionality is a key ingredient in an explanation of how language works.

---

<sup>8</sup> Of course there are ambiguities that are obviously structural, like the one in *Old men and women were asked to get in the lifeboat*. The assumption here is just that *some* non-lexical ambiguities are not structural, with the different scopes of (1) and (2) in section 3 as paradigmatic examples.

## 8. Compositionality in a relational framework

To answer the question just posed it is necessary to separate two formulations of compositionality, which are (almost) indistinguishable in the single-valued case. I call them the *Rule* version and the *Substitution* version, and only give approximate and non-technical formulations here; this suffices to illustrate the main points. First, consider the single-valued case:

(Rule) For every syntactic rule  $r_{\text{syn}}$  there is a corresponding semantic operation  $r_{\text{sem}}$  such that the meaning of the result of applying  $r_{\text{syn}}$  to certain expressions is the result of applying  $r_{\text{sem}}$  to the meanings of those expressions.

That is, the meaning of the whole is determined by the meanings of the parts, and the mode of composition ( $r_{\text{syn}}$ ).

(Subst) Replacement of (not necessarily immediate) sub-expressions of a complex expression by synonymous expressions results – if the result is meaningful at all – in a complex expression synonymous with the original one.

That is, replacement of synonymous parts (when allowed) preserves meaning. It is not hard to give precise versions of (Rule) and (Subst), and to show that, provided sub-expressions of meaningful expressions are always meaningful (which (Rule) but not (Subst) presupposes; cf. footnote 3), they are equivalent (see Hodges 2001 or Westerståhl 2004).

(Rule) presupposes a single-valued meaning assignment: witness phrases like “*the* meaning of”. (Subst) does so too, in a less perspicuous way, in that it uses a straightforward notion of synonymy: two expressions are synonymous if (both are meaningful and) they have *the same* meaning.

Clearly it is the (Rule) formulation that best corresponds to the intuitive motivation behind compositionality: it offers an explanation of linguistic productivity in terms of knowledge of syntactic rules and their semantic counterparts. Roughly, the first explains why we recognize new sentences as well-formed, and the second why we effortlessly see what they mean.<sup>9</sup> And it is this formulation that readily extends to the case of a relational semantics, where an expression may have several meanings:

(Rule<sub>Amb</sub>) For every syntactic rule  $r_{\text{syn}}$  there are corresponding semantic operations  $r_{1,\text{sem}}, \dots, r_{k,\text{sem}}$  such that the meanings of the result of applying  $r_{\text{syn}}$  to certain expressions is the result of applying  $r_{i,\text{sem}}$ , for some  $i$ , to meanings of those expressions.

This formulation is still imprecise, and there are in fact some niceties pertaining to finding an exact version of (Rule<sub>Amb</sub>). I will not go into them here. But I think the idea is clear: For each syntactic rule there is now (possibly) more than one corresponding semantic rule. Thus, finding a meaning of the complex expression involves having access to some meanings of its immediate constituents, and applying one of the corresponding semantic operations to those meanings. One must choose, as it were, meanings of the parts, and also the semantic operation to apply.

---

<sup>9</sup> Of course, much more is required for such an explanation to work, in particular that the semantic operations are *easy* (for us) *to compute*, but that is another issue.



This is a little more complex than in the case when there are unique meanings of the parts and just one operation to apply, but not much. After all, already with lexical ambiguity one has to choose among the possible meanings of lexical items; the only new element here is that one also has to choose a rule. But this seems equally feasible; all the available rules are specified in advance, and depend only on the syntactic ‘mode of combination’, just as before.

Consider our sample sentences (1) and (2) in section 3. The syntactic structure of both is the tree  $T$ . With a relational semantics, and assuming there are no lexical ambiguities, each sentence has two meanings, corresponding to the two possible ‘orders’ between the quantifiers. Suppose ambiguity sets in only at the last rule application, the one that forms the sentence from the NP and the VP (other accounts are also possible). At this stage there are *two* semantic rules: one for each scoping. We saw that for these two sentences, the natural reading involves one scoping for (1) and the other one for (2). This choice can clearly be described as a choice between one of two semantic rules, both of which are given, and ‘known’ by the speaker, in advance.

I end with three comments, stated here without proof. They can be proved rigorously in a suitably precise setting.

*First*, the precise version of  $(\text{Rule}_{\text{Amb}})$  is such that, when the relational semantics happens to be single-valued,  $(\text{Rule}_{\text{Amb}})$  is equivalent to the usual  $(\text{Rule})$ .

*Second*, even when  $n = 1$  in  $(\text{Rule}_{\text{Amb}})$ , so that each syntactic rule corresponds to exactly one semantic operation, there is still room for lexical ambiguity. One can show that  $(\text{Rule}_{\text{Amb}})$  in this case gives a reasonable account of how lexical ambiguity is propagated to complex expressions. More generally,  $(\text{Rule}_{\text{Amb}})$  is intended to work in the presence of all three kinds of ambiguity: lexical, structural, and the third kind exemplified in section 3.

*Third*, a difference with the single-valued case, however, is that there seems to be no obvious substitution version of  $(\text{Rule}_{\text{Amb}})$ . The reason, I think, is that there is no obvious relation of synonymy in relational semantics. There are several natural synonymy relations: the strongest one saying that two expressions are synonymous if they have exactly the same meanings. (This is the synonymy corresponding to the set semantics of section 6.) But there are also natural weaker versions. It can be shown that none of these is equivalent to  $(\text{Rule}_{\text{Amb}})$ . Apparently, with a relational semantics, the notion of synonymy becomes more problematic.

## **9. Conclusion**

I have argued that even if we give up the idea that each (structured) expression has a unique meaning, the idea of compositionality remains, and so does the explanation of the ‘remarkable facts’ about human language that Frege and many others have noted. The main part of the argument was an indication of how to formulate the principle of compositionality within a relational semantics, a formulation that extends the standard version for functional semantics. Clearly, for the argument to go through one needs a

precise version of this principle, and an investigation of its properties, just as the usual version has been investigated and is now quite well understood. Also, one needs to look at a variety of applications to concrete cases. These are things I hope to come back to in the future.

## References

- Barker, C. and Jacobson, P. (eds.), *Direct Compositionality*, Oxford University Press, Oxford, 2007.
- Cooper, R., *Quantification and Syntactic Structure*, D. Reidel, Dordrecht, 1983.
- Debusmann, R., Duchier, D., Koller, A., Kuhlmann, M., Smolka, G., and Thater, S., A relational syntax-semantics interface based on dependency grammar, in *Proceedings of the 20th International Conference on Computational Linguistics (COLING)*, Geneva, 2004, 176–182.
- Dowty, D., Compositionality as an empirical problem, in Barker and Jacobson (2007).
- Hodges, W., Formal features of compositionality, *Journal of Logic, Language and Information* 10 (2001), 7–28.
- Pagin, P. and Pelletier, F. J., Content, context and composition, to appear in G. Peter and G. Preyer (eds.), *Content and Context. Essays on Semantics and Pragmatics*, Oxford University Press, Oxford.
- Pagin, P. and Westerståhl, D., Compositionality, to appear in K. von Stechow et al. (eds.), *Semantics: An International Handbook of Natural Language Meaning*, Mouton de Gruyter, Berlin.
- Pelletier, F. J., Semantic compositionality: Free algebras and the argument from ambiguity, in M. Faller, S. Kaufmann, and M. Pauly (eds.),

*Proceedings of Seventh CSLI Workshop on Logic, Language and Computation*, CSLI Publications, Stanford, 1999.

Szabó, Z. G., Compositionality, *The Stanford Encyclopedia of Philosophy* (Spring 2007 Edition), E. N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2007/entries/compositionality/>.

van Deemter, K., Towards a logic of ambiguous expressions, in K. van Deemter and S. Peters (eds.), *Semantic Ambiguity and Underspecification*, CSLI Publications, Stanford, 1996, 203–237.

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