Semantic Compositionality:

Free Algebras and the Argument from Ambiguity

Francis Jeffry Pelletier

Semantic Compositionality is the principle that the meaning of a syntactically complex expression is a function (only) of the meanings of its syntactic components together with their syntactic mode of combination. Various scholars have argued against this Principle, including the present author in earlier works. One of these arguments was the Argument from Ambiguity, which will be of concern in the present article. Opposed to the considerations raised against the Principle are certain "formal" arguments that purport to show that there is no empirical content to the Principle. One of these formal arguments makes use of the notion of *free algebras*. The present article investigates the relationship between these two types of argument.

1 Introduction and Background

Some of the more informal writings on Semantic Compositionality — the principle that the meaning of a syntactically complex expression is a function (only) of the meanings of its parts plus their syntactic mode of combination — have alleged that Compositionality is "incompatible with certain obvious facts concerning meaning and inference" ([Pell94a]). One category of such alleged facts concerned synonymy (for which see [Pell94b]) while another category concerned ambiguity.

Despite the informal arguments of the sort given in the papers just mentioned (and in papers by many other authors), there is also a tradition of papers which allege that Semantic Compositionality "is not an empirical issue." The reasoning for this claim can be seen from a representative sampling of quotations:

Proceedings of the Seventh CSLI Workshop on Logic, Language and Computation.
Martina Faller, Stefan Kaufmann & Marc Pauly (eds.)

> ...any semantics can be encoded as a compositional semantics, which means that, essentially, the standard definition of compositionally is formally vacuous. [Zad94, p. 329]
>
> Janssen's argument shows that any recursively enumerable language...can be assigned any set of meanings in a compositional way. It follows that compositionality is not an empirical principle, but a methodological one. [Hen93, p. 137]
>
> ...given any connection of syntactic operations with semantic operations (of the same number of arguments), any arbitrary map from basic lexical items to suitable semantic entities is uniquely extendible to a homomorphism. Thus...*by itself, compositionality provides no significant constraint on semantic theory*. [vBen84, p. 57, italics in original]

The thought in these author's minds is this: they believe (i) it can be proved that "every semantics can be encoded as a compositional semantics", and (ii) it follows from this that "compositionality provides no significant constraint on semantic theory" and hence that "compositionality is formally vacuous."

Although these two claims have been discussed, and debunked, in [West98]; there seems nonetheless to be some use in looking a little more closely at the issues involved. We start by looking at the informal "Argument from Ambiguity."

2 The Argument from Ambiguity

The intuitive picture behind most accounts of semantics, both the formal and the informal, is that the syntactically-analyzed expressions of natural language are mapped onto some "meaning structure" that provides an "interpretation" of the expression. Different theories of semantics will of course have different conceptions of what this meaning structure is and therefore of what an interpretation is. But at this very high level of generality we

can embrace pretty much every conceptualization of semantics — from the most symbolically formal to the psychological and social accounts of meaning and to the very informal views that often find expression in related literature. For it seems that any semantic theory which holds its role to be a matter of interpreting a symbol system would be encompassed by such a general account, and I think that all such theories are subject to the following considerations concerning ambiguity.[1]

Most instances of this wide variety of semantic theories acknowledge that there is ambiguity — that some expressions of the symbol system can be interpreted in more than one manner. But since compositionality requires that the interpretation process be *functional* [that each item of the symbol system is to be associated with no more than one meaning], it might naturally concluded that a compositional semantics cannot admit ambiguity. But theorists who advocate a compositional semantics respond by pointing to not just one, but *two* ways to allow ambiguity while maintaining compositionality.

The first is to allow lexical items to be (primitively) ambiguous; that is, we admit that individual words can be ambiguous without violating the spirit of compositionality. Then, not only is an individual word allowed to be ambiguous, but also complex phrases containing the word can be ambiguous because of the presence of that word. For example, it might be claimed

(1) She sat beside the bank

is ambiguous due to a lexical ambiguity on the part of *bank*. 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 meanings 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.

The second way to accommodate ambiguity is to notice that compositionality is only defined on syntactically-analyzed struc-

tures. (It is in this way that the "manner of syntactic combination" clause in the definition of 'compositionality' is incorporated). But it might turn out that two different ones of these syntactically-analyzed structures could employ the same lexical items in the same order, despite the two structures being different. For example, one might claim that

(2) Flying airplanes can be dangerous

is such a sentence — that there are two distinct syntactic structures here that happen to embody the same words in the same order. When faced with such cases we might agree to call this a kind of ambiguity. However, we should be aware that this type of ambiguity is not a challenge to compositionality because of the two syntactic structures involved, indicating that there were two different "manners of combination." And of course, with the two different modes of combination we should expect two different meanings for the two structures.[2]

But 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. Are there any such structures? [Pell94a] suggests the following:

(3) a. Every linguist knows two languages

b. John wondered when Alice said she would leave

c. When Alice rode a bicycle, she went to school

d. The philosophers lifted the piano

e. The Canadian family used less water last year than the preceding year

First, he argued that these are each semantically ambiguous. (3a) exhibits the well-known distinction between there being two languages (e.g., English and French) that are known by every linguist (on the one hand) and that each linguist knows at least two languages — which may be different from those known by

other linguists (on the other). (3b) displays a difference in what John is wondering: is it when Alice is leaving? or when she said something? (3d) is ambiguous between (i) stating that, (in the past) on those occasions when Alice rode her bicycle, she used it as a means of transportation to go to school, and (ii) in those times of yore, back when Alice used to ride bicycles, she was a student. (3d) manifests the familiar distributive/collective distinction: the philosophers each lifted the piano vs. they acted as a group to produce a lifting of the piano. And (3e) is ambiguous between whether we are talking about the sum total of all the water used by all Canadian families vs. whether we are talking about what the average Canadian family uses. (There are circumstances where one is true and the other false, as for example when the number of Canadian families changes from one year to the next.)

In addition to arguing that the examples in (3) are semantically ambiguous, [Pell94a] also argues that each of them has but one syntactic analysis, and that the only reason anyone would think otherwise is that she or he was antecedently committed to compositionality. For example, although we happily represent the two *meanings* of the first example by different scopes of quantifiers, it is extremely difficult to see that *in the syntax* there is any such ambiguity. For all the world it appears that there is exactly one phrase structure for that sentence. It is especially hard to see that such (quasi-) syntactic mechanisms as Montagovian "quantifying in" rules ([Mont73]) or "Cooper storage" ([Coop83]) are motivated by anything more than a desire to accommodate compositionality. Isn't it clear that no one not already in the grips of a theory would say that there is more than one syntactic structure in the first example? Similarly, isn't it hard to convince oneself that there are two different structures in the second example ...one with a "gap" in a location that the other one doesn't have? Surely that's pushing the issue more than the evidence warrants. And where will one find any structural ambiguity in the third example? Perhaps the compositionalist will be forced to try to locate an ambiguity in *when*? Sometimes it "joins" the two clauses as a means–ends but other times it merely says that

two types of events occur during the same interval? Doesn't all this seem pretty far-fetched?[3]

Although [Pell94a] argues vigorously that this Argument from Ambiguity shows that compositionality is to be given up, there is a sense in which he does not think it is conclusive. His view is that whether there should be compositionality in semantic theory is an empirical issue; he just believes that the best explanation of the ambiguity facts is that compositionality is false. He in fact canvasses various strategies that could be used to maintain compositionality in semantic theory (such as the "quantifying in" rules and the "Cooper storage" strategy, as well as some other gambits), admitting that if they were adopted then this class of sentences would be given a compositional semantics. As mentioned, it is his view that the burden of doing these "unnatural" maneuvers outweighs whatever solace comes from having compositionality. But he views it as an empirical matter in much the same way as is adopting any highly general, theoretical position in a science. And many of the same types of considerations come into play in evaluating whether compositionality should be adopted in semantic theory as whether any specific abstract principle should be adopted in some science. So in the end it might turn out that the best overall account of the data is achieved by adopting compositionality. In this "best overall account" it may turn out that the sentences in (3) *are* syntactically ambiguous, or perhaps that they are *not* semantically ambiguous.

The upshot of his discussion is that the issue of ambiguity is empirical in nature. The decision about whether one should accept these claims, both about the alleged semantic facts (that the sentences in (3) are ambiguous) and the alleged syntactic facts (that there is but one syntactic structure in each one of them), is to be decided in whatever manner one in general discovers linguistic facts. If this general method is a matter of introspection of one's linguistic abilities, then so is the discovery of the truth concerning ambiguity; if the general method is a matter of canvassing the usage of populations, then so is the discovery of truths concerning ambiguity; if the general method is more theoretical (e.g., discovering how certain possible positions fit into

a broader theory), then so is the discovery of truths concerning ambiguity. Most probably, empirical linguistic investigation embodies all three of these considerations, and others also; and so, the discovery of the truth concerning the ambiguity facts will also. But it is to be emphasized that these facts, whatever they may turn out to be, are empirical in nature. It is *not* a matter of a theorist deciding that "there's more work to be done in the construction of an underlying syntactic theory so that there will be no ambiguity", even if it is possible for a syntactic theory without ambiguity to be constructed (and still generate all the sentences of the language). For, the existence or non-existence of ambiguity is an *empirical* matter—not something that can simply be dictated by aesthetic taste.

But if that's the assessment of the Argument from Ambiguity, what is to be made of the claims mentioned at the outset of this paper alleging that Compositionality is "formally vacuous" or "provides no significant constraint" and "is not an empirical issue"? These claims are backed up by proofs, and I would like to investigate how one of them (at least) interacts with the Argument from Ambiguity.

3 The Argument from Free Algebras

The mathematical results mentioned or alluded to in the quotations at the beginning of this paper are at odds with the conclusions just cited from [Pell94a]. According to Pelletier's conclusions, it is a matter of empirical theory-construction whether to honor the appearance of semantic ambiguity and the non-appearance of syntactic ambiguity in examples like those cited in (3), or instead to postulate that they aren't ("really") semantically ambiguous or that there ("really") is some syntactic ambiguity present. But according to the authors of those quotations this is not an empirical issue at all, because there is a compositional semantics for these examples.

I wish to turn to an examination of one of these formal arguments, the one alluded to in this slight expansion of van Benthem's quotation as given above.

> ...the syntactic algebra ...is *free* (it is freely generated by the basic lexical elements). What this algebraic assertion amounts to is this. ...[G]iven any connection of syntactic operations with semantic operations ..., any arbitrary map from basic lexical items to suitable semantic entities is uniquely extendible to a homomorphism.

van Benthem is here employing an algebraic interpretation of syntax and semantics, so as to be able to bring results from universal algebra to bear on the topic. There is nothing to take exception to in doing this, so long as the features of the objects of study (the linguistic syntax and the semantics) are faithfully mirrored in the particular algebras being employed.

As the expanded quotation shows, van Benthem assumes that the syntactic algebra is *free*, that is, that it is freely generated from its lexical items. What does this mean? [End72, p. 27] defines

> [A set] C is *freely* generated from [another set] B by functions f and g iff in addition to ...being generated, we have
>
> 1. f_C and g_C are one-to-one, and
>
> 2. The range of f_C, the range of g_C, and the set B are pairwise disjoint.
>
> (Here f_C and g_C are the restrictions of f and g to [the set] C.)

And using this definition he states his "Recursion Theorem", giving it the following informal gloss

> Assume that the subset C of U is freely generated from B by f and g .... Viewed algebraically, the conclusion of this theorem says that any map h of B into [some set] V can be extended to a homomorphism h from C into V.

It is perhaps easiest to see the import of all this by example. Consider the sentence logic with only two connectives, $\neg$ and $\rightarrow$; this will be our set C. In the definition of 'freely generated' f_C and g_C correspond to the syntactical operations of forming complex sentences from simpler ones by putting them together with $\neg$ and $\rightarrow$. The set B is the set of atomic sentence letters. Notice that f_C and g_C are both one-to-one: forming complex sentences by these operations yields a unique complex, and furthermore each complex sentence is uniquely decomposable into its simpler components. We also note that the ranges of f_C and g_C are disjoint (negations and conditionals are disjoint), and that furthermore these are distinct from the sentence letters (sentence letters are disjoint from negations and from conditionals). So therefore, the set of sentences is freely generated from the set of sentence letters by the operations of negation and conditionalization. Thus the Recursion Theorem applies to this language: any map of the sentence letters into any set V (e.g., into meanings) can be extended to a homomorphism from all sentences into V. And therefore the sentence logic (with $\neg$ and $\rightarrow$) has a compositional semantics.

One can attack the deployment of this argument in two ways. The first is taken by [West98] who correctly remarks that although we will be guaranteed *some* compositional semantics for this language it is not the case that it will (necessarily) correspond to our pre-theoretical ideas of the appropriate semantic values for our sentences and other complex terms. I will not follow up this line of attack on the use of the notion of free generation here. Instead I turn to a second direction of deflecting the use of free generation.

The crucial assumption that has been made here is that the entire language can be freely generated from the sentence letters. Well, of course this is true in the case of the sentence logic: each sentence is uniquely decomposable into its immediate constituents. There is never a case where a sentence made up in one manner (e.g., by $\neg$) is identical with a sentence made up in a different manner (e.g., by $\rightarrow$).

But we should wonder if this is an appropriate assumption

for natural language. Note that the idea of freely generating a language from its lexicon is exactly to deny ambiguity of the sort outlined above, where there is no ambiguous lexical item, where there is only one syntactic analysis of the sentence, and yet where there are two meanings to be associated with that syntactically analyzed form. Isn't it really difficult to believe that English (or any other natural language) is freely generated from its lexical items? That there can never be cases where precisely the same syntactically-defined sentence can be generated from two different analyses? It doesn't happen in the sentence logic case because there are devices such as parentheses involved in constructing the complexes, otherwise strings such as $\neg P \rightarrow Q$ would show that even sentence logic is not freely generated from the atomic sentences. But in ordinary language these devices simply are not in play. There is a very large number of syntactic rules in play which can interact in complex ways. And it seems hard to find *any* justification for saying that all complex structures are freely generated from the lexical items.

Or at least, if one insists that natural languages are freely generated by their lexicon, then that is just another way to insist that natural languages are compositional. And this would be *not* to treat the issue of ambiguity as empirical, but instead to decree *a priori* that there cannot be any since languages must be compositional. Not a very satisfying position, I should think.

Similar considerations carry over to the theory of [Mont70], where the background disambiguated language *is* freely generated, and therefore it can be given a compositional semantics. But it is clear that the claim of compositionality either does not carry over to natural language, or if it does carry over (in some extended sense) then this extended sense will beg the question of ambiguity exactly as much as the accounts just described. The former alternative, according to which compositionality does not carry over to natural language, follows the letter of Montague's picture in which the strings of the natural language are not ascribed *any* syntactic analysis at all. In such a framework the issue of compositionality does not arise at all for natural language. All we can say is that the background disambiguated language is

compositional, but that this has no implications for natural language. The second alternative, admittedly a non-literal interpretation of Montague's picture, would understand the "association" of a set of members of the disambiguated language with a sentence of natural language as a way of saying that the sentence has the analysis (or analyses, if the set has many members) specified by the members of this "associated" set. With this interpretation the issue of compositionality *does* come up, but it is to be understood in the same way as the theories discussed in the previous paragraph. In this view, each of the ambiguities is attributed to a distinct syntactic analysis of the background disambiguated language. So, all ambiguities are, once again, to be traced to lexical items or to syntactic differences. But as before, this view merely presumes the falsehood of (what I've claimed to be) the empirical question: Is there any non-lexical ambiguity which is not to be traced to the operation of some syntactic rule? The use of a disambiguated language that is freely generated from its lexical items does not *prove* that "any semantics can be converted to a compositional semantics," but instead *presumes* it.

4 A Few Concluding Remarks

I have focused in this paper on the issue of ambiguity and what light it can throw on the topic of semantic compositionality. I argued that it is an empirical issue as to whether there is ambiguity of the sort alleged for the examples of (3), where the ambiguity is not to be traced to an ambiguous lexical item nor to a surface structural distinction. If there is such a type of ambiguity — and I argued that there was — then compositionality is shown to be empirically false. Of course, since I think this is an empirical issue, I am bound to concede that perhaps the best account of the facts related to ambiguity will show that there is no such thing as an ambiguity of the sort alleged for (3). In that case compositionality will have survived this challenge.

But in either of those cases the issue is to be decided empirically. It cannot be decreed from the "outside" by importing considerations that do not genuinely apply to the case at hand.

I looked at one particular form of this illegitimate importation, that of the notion of having a language be freely generated from its lexicon. I argued that at best this is a form of begging the question, since free generation denies ambiguity. An important moral to draw from this is to notice: any mathematical proof about matters in this realm will have to start with certain empirical assumptions, or else its conclusion will have no empirical content. But then it behooves us to investigate the empirical support the initial assumptions have before we embrace the conclusion.

I think much the same point could be made against the use of [Jan86]'s argument and [Zad94]'s argument to show that compositionality is not an empirical issue. But I shall not undertake that here, since they are not explicitly engaged with the Argument from Ambiguity with which I started. (See [West98] for further comments).

The fact is, I believe, that none of the "formal arguments" succeed in showing that "compositionality is vacuous." And this is because they all make unjustified assumptions about what the syntactic and semantic facts are allowed to be for a natural language. It would be a very interesting task—and one that I hope to undertake sometime in the near future—to describe exactly what one would *really* have to prove if one were to convincingly demonstrate that "every semantics can be converted into a compositional semantics."

5 Acknowledgments

The impetus for this paper derives from many sources. Most immediately was Johan van Benthem and Martina Faller who wanted "something on compositionality" to be in the Proceedings of the Seventh Logic, Language and Computation conference. But this in turn was predicated on work that I was pushed into doing by David Beaver at meetings of his Semantics and Pragmatics workshop. I thank the many participants of these meetings, especially Frank Veltman, David Beaver, and Stanley Peters. Of course my earlier writings on compositionality

have aroused many angry (or dismissive) conversations with various scholars, both in the form of electronic spitballs and in the form of (metaphorical) pies in the face when I wasn't looking. I would especially like to thank Wilfrid Hodges, David Israel, Theo Janssen, Bernard Linsky, Ivan Sag, and Dag Westerståhl. Further thanks go to Dag for allowing me to see a preprint of [West98], which has influenced my thinking on this topic a lot. Two anonymous referees also made very useful comments, which I hope I have addressed in this version of the paper. Finally, much of the work on this paper was done while I was visiting CSLI at Stanford University, and I would like to thank John Perry and Ed Zalta of CSLI for providing me a very comfortable place indeed to do research, and for not throwing any pies in my face. And most importantly I want to thank the University of Alberta, which awarded me a McCalla Professorship for the 1998-99 academic year, freeing me from teaching so as to do research. This help is most gratefully acknowledged.

Notes

[1]Well, perhaps this characterization doesn't quite capture the following conception: Sentences of natural language are *strings* that have no grammatical structure "in nature." Instead, we theorists devise another, "formal" language which contains syntactic structures and categories, and we use this as a "model" for the natural language by associating any natural language sentence with a particular set of items in the model. For example, we might decide to associate every item in the model that has the same terminal symbols in a certain order with the natural language sentence that also has those terminal symbols in that order. If there is but one such element of the model, then it alone is associated with the natural language sentence; but if there are ten such elements, then the entire set of ten is associated with the natural language sentence. But it is important to note that *no* natural language sentence is claimed directly to have a syntactic structure: there is just an "association" of elements of the model with natural language items. (This is the sort of attitude taken by [Mont70]). So strictly speaking, this conception of the relation between natural language and a model language that describes structure falls outside the class of theories I've introduced here. Still, I believe my later remarks will apply even to this conception; and I will return to it shortly.

[2]The view of natural language analysis attributed to [Mont70] in the last footnote can be seen to yield an account of ambiguity quite similar to the account given in this "second type." The difference is that in this second

type, the natural language sentence itself is given multiple analyses, whereas according to Montague's strategy the different elements of the disambiguated language are merely "associated with" the natural language sentence and Montague does not further characterize this relationship. Nevertheless, it seems to me that with a suitable generosity of interpretation, it would not be unfair to categorize his conception as being cut from the same cloth as my "second type", at least so far as understanding what syntactic ambiguity amounts to.

[3]In the conception of [Mont70]], the alleged ambiguities found in these sentences would be attributed instead to different "analysis trees" in the underlying disambiguated language. The idea would be that the natural language sentences listed in (3) would each be "associated with" a group of more than one analysis trees. Again, rather than attributing structural ambiguity to the natural language sentences directly, this view could find the ambiguity in the fact that there are many different items in the disambiguated language that are to be "associated with" the natural language sentence.

As with the view that the natural language sentence itself is syntactically ambiguous, it seems clear that this view is also not a challenge to compositionality, because it finds the ambiguity to reside in the existence of multiple members of the background disambiguated language being "associated with" the same natural language sentence.

References

[vBen84] van Benthem, J. (1984) "The Logic of Semantics" in F. Landman & F. Veltman (eds.) *Varieties of Formal Semantics* (Dordrecht: Foris) 55–80.

[Coop83] Cooper, R. (1983) *Quantification and Syntactic Theory* (Dordrecht: Reidel).

[End72] Enderton, H. (1972) *A Mathematical Introduction to Logic* (NY: Academic Press).

[Hen93] Hendriks, H. (1993) *Studied Flexibility. Categories and Types in Syntax and Semantics*. Dissertation, University of Amsterdam.

[Jan86] Janssen, T. (1986) *Foundations and Applications of Montague Grammar*. CWI Tract #19, Centre for Mathematics and Computer Science, University of Amsterdam.

[Mont70] Montague, R. (1970) "Universal Grammar" in R. Thomason (ed.) *Formal Philosophy* (New Haven: Yale Univ. Press), pp. 222–246.

[Mont73] Montague, R. (1973) "The Proper Treatment of Quantifiers in Ordinary English" in R. Thomason (ed.) *Formal Philosophy* (New Haven: Yale Univ. Press), pp. 247–270.

[Pell94a] Pelletier, F.J. (1994a) "The Principle of Semantic Compositionality" *Topoi* **13**, 11–24.

[Pell94b] Pelletier, F.J. (1994b) "Semantic Compositionality: The Argument from Synonymy" in R. Casati, B. Smith, G. White (eds.) *Philosophy and the Cognitive Sciences* (Vienna: Hölder-Pichler-Tempsky) pp. 208–214.

[West98] Westerståhl, D.(1998) "Purported Mathematical Arguments for the Vacuity of Compositionality" *Linguistics and Philosophy* **21**, 635–643.

[Zad94] Zadrozny, W. (1994) "From Compositional to Systematic Semantics" *Linguistics and Philosophy* **17**, 329–342.