

On the Possibility of Compositional Pragmatics

Peter Lasnik
University of Illinois

1. Introduction. The question we were asked to consider for this workshop is whether contextualism — understood as the claim that context-sensitivity goes beyond indexicality and affects every natural language sentence — is incompatible with compositionality, and threatens the project of building a systematic semantics for natural language.

There are two ways of taking this question. One is to focus on the issue of compatibility: Assuming context-sensitivity is pervasive, what does that mean for compositionality? The other is to focus on this issue of whether contextualism is well supported: Assuming that contextualism is incompatible with compositionality, are the arguments for it strong enough to pose a real threat? In this paper I will take the first approach, set aside the issue of just how context-sensitive interpretation really is, and consider the question of whether there is an incompatibility between pervasive context-sensitivity and compositional interpretation in the first place.

I will argue that there isn't any such incompatibility, at least if we understand compositionality in one fairly standard way, namely as a principle that interpretation be assigned homomorphically, and construe this as a constraint on structural patterns in the language, rather than a model of what mental procedures people actually perform as they are interpreting each other's utterances. More specifically:

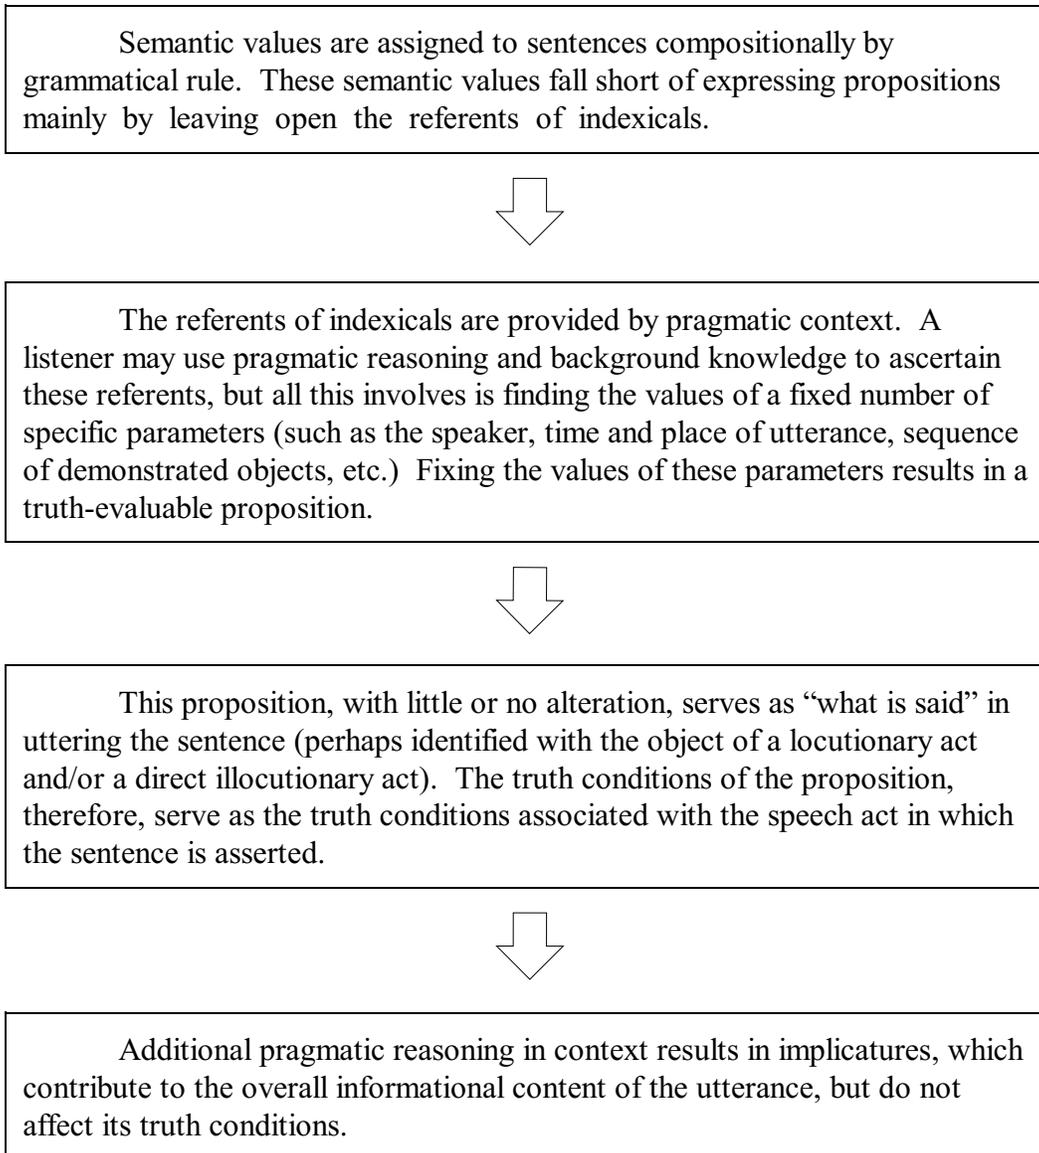
- We should carefully distinguish between the question of how a *grammar* assigns interpretations, and the question of how *people* figure out what interpretation the grammar assigns;
- The fact that people must rely on pragmatics in the latter task does not show that interpretation is assigned non-compositionally;
- The claim that interpretation is assigned homomorphically is equivalent to a familiar substitutivity principle — implying that counterexamples must have a certain form, and that arguments against compositionality must address the intuition underlying the substitutivity principle;
- Thoroughgoing contextual effects on truth conditions can be handled compositionally in a grammar if that grammar is allowed to *invoke* pragmatic concepts; this does not mean that the grammar must *explain* pragmatic concepts, or do all the work of a pragmatic theory.

2. Contextual effects and the semantics/pragmatics “hand-off.” Why would we think that pervasive context-sensitivity threatens compositionality in the first place? I actually don't know of any arguments that have been given specifically and explicitly just against compositionality based on contextualism. The arguments that I am familiar with are not so tightly focused against

compositionality per se, but are really advanced against a much broader conception of how semantics, truth conditions, and pragmatics relate to each other — a conception that includes compositionality as a fairly central principle, but which also involves a number of other assumptions about linguistic architecture and how it gets deployed in actual talking and understanding.

Often, this conception is attributed to Grice, or to followers of Grice, and can be summarized in the following chart:

(1)



This looks like a pretty simple picture, but there is actually a lot packed into it, and a lot of different points at which one might try to attack it. What I take to be the “classical” attack from contextualism goes something like this:

(2)

Exhibit a sentence whose truth conditions seem to depend on context.
(E.g. Travis: *The leaves are green*, said of painted leaves — intuitively true in a context where a photographer is setting up a shot, false in context where a biologist is seeking a species with naturally green leaves.)



Argue that this contextual difference in truth conditions is not plausibly the result of fixing the value of some definite parameter, as in indexical resolution.



Infer that a truth-evaluable proposition cannot be obtained by fixing values for a limited set of contextual parameters in semantic values assigned to sentences compositionally by grammatical rule.



Show that the variation in truth conditions can be explained by appeal to the practical purpose of the utterance and/or pragmatic, non-grammatical reasoning on the part of interpreters.



Conclude that truth conditions are significantly underdetermined by grammar, but must be assigned in part by “global,” non-compositional, pragmatic processes.

Several people have given this kind of argument. My presentation of it here is perhaps closest to the version in Travis (1997), but roughly similar arguments can be found in Carston (1988), Recanati (2005) and many other places.

I'd like to point out that this argument is really primarily about the semantics/pragmatics interface, and just how close to the truth conditions of an utterance we can get using just the grammatical system of a language. It is about compositionality only under the assumption that grammatically-based interpretation is compositional, and pragmatically-based interpretation is not.

Of course, that is a pretty widespread assumption, and it would seem to be true almost as a matter of definition. After all, compositional interpretation is directly sensitive to the division of an expression into constituent parts; that division is something given by the grammar of the language, and so compositional interpretation is grammatically sensitive. Pragmatics, on the other hand, is extra-grammatical, and therefore can't be expected to be as tightly tied to grammatical description as it would have to be if it were restricted by compositionality.

I will argue that we should *not* assume that all pragmatic reasoning is unconstrained by compositionality (though of course it will depend in part on what, exactly, we mean by “pragmatic”). In particular, when pragmatics is used to figure out what the grammar gives us, it can be expected to respect compositional assignment by the grammar.

As a first step toward clarifying this idea, let us return to the Gricean or quasi-Gricean picture outlined in (1), which was the target of the contextualist attack. I would like to point out that there is something odd about this picture, completely independent of any considerations having to do with contextual variation in truth conditions.

Specifically, the diagram in (1) is easily construed as depicting an interpretive *procedure*, in which something is produced by the semantics (construed as a component of grammar) and then “handed off” to pragmatics, which is not.

Semanticists talk this way all the time, and most of the time it does not harm. But if we take this conception of interpretation very seriously, it seems to me that we run the risk of making a conceptual error in talking about grammar “handing off” anything to pragmatics.

With some exceptions, of course, I think most semanticists would take the position that a grammar (including the semantic component of a grammar) is not supposed to model what people do, on-line, in real time, as they are speaking and listening. For Chomskyans, a grammar is supposed to be a model of competence, not performance. Others might not even go that far, but would simply say that the grammar is a model of patterns in language, and leave open the issue of just how directly those patterns are encoded in our mental representations, even at the level of competence. Either way, a grammatical theory, including the semantic component of a grammar, is not supposed to be a direct description of what we do as we are engaged in interpreting each other's utterances.

I suspect, therefore, that not many semanticists would claim that what we do when we *are* engaged in this task is to work through a formal derivation defined directly by the syntactic rules of our grammar, and through a compositional truth definition defined directly by the semantic rules. And if we don't do that, there wouldn't seem to be any point in the actual interpretation process at which the grammar “spits something out” to be handed off to pragmatics — at least if we understand pragmatics to be about what people actually do in the act of interpreting each other's utterances. Trying to locate the output of a grammar at some point in the process that people engage in as actual linguistic performance — whether production or interpretation — is to make a category mistake; it is to misconstrue what a grammar is a theory of.

Now, perhaps not everyone does understand pragmatics as a theory of what people actually do. In reading Grice, for example, I don't get a clear sense that he intended his theory of conversational implicature as a theory of the actual psychological processes that people go through in real time, as opposed to a theory of what the rational justification is for that kind of inferencing, however it might be implemented psychologically. But certainly, in a great deal of theoretical discussion of pragmatics, the assumption seems to be strong that the inferential procedures that a pragmatic analysis appeals to are supposed to be procedures that people actually engage in during the act of interpretation. If that is what is intended, we should not think of the grammar as "handing anything off" to those procedures; that would make sense only if people worked through derivations defined by the grammar as part of real-time interpretation.

3. *Interpreting without computing a derivation defined by the grammar.* No doubt some theorists will want to deny that grammatical knowledge is anything other than knowledge how to engage in linguistic performance — in which case a grammatical theory *would* model what we do in real-time processing (and it would make sense to consider whether a "hand-off" from semantics to pragmatics takes place in that processing). But the following simple example will illustrate that at least in some cases, what we do psychologically in interpreting a sentence couldn't *possibly* be computing the operations defined directly by the semantic component of the grammar. This example is discussed in more detail in Lasnik (2009). It uses a simple artificial language, but the artificiality of the example should not undermine the theoretical point.

Let's call this language L . The vocabulary of L includes the vertical stroke $|$, the cross-hatch symbol $\#$, and a predicate \mathbf{TL} . The syntax has the following formation rules:

(3) *Syntax of L :*

- $|$ is a number term
- If α is a number term and β is a number term, then $\alpha\beta$ (the concatenation of α with β) is a number term
- If α is a number term, then $\#\alpha$ is a string term
- If α is a string term, then $\mathbf{TL}(\alpha)$ is a sentence

Semantically, we let number terms denote numbers:

(4) *Semantics of number terms:*

- $\llbracket | \rrbracket = 1$
- If α is a number term and β is a number term, then $\llbracket \alpha\beta \rrbracket = \llbracket \alpha \rrbracket + \llbracket \beta \rrbracket$

So, for example, $\llbracket ||| \rrbracket = 3$.

String terms will denote strings in a second language, L_{NR} . We require that L_{NR} be a non-recursively-enumerable language over a 26-letter alphabet. Although L_{NR} is non-recursively-enumerable, we make the standard assumption that it is a set of finite strings over a finite alphabet, hence that it is only countably infinite in cardinality. So, there will be a function putting it in 1–1 correspondence with the set of natural numbers \mathbb{N} . In fact, there will be many such functions, including one which arranges the strings of L_{NR} in order of increasing length. Call that function f . (We may stipulate that f arranges strings of the same length in alphabetical order, or in any other convenient way.)

Now, even though f is in some sense systematic — it sorts strings by length — it is not a computable function. By a *computable function* we just mean one which can be represented as a recursively enumerable set of ordered pairs. But if we could recursively enumerate the pairs in f , we could easily construct a recursive enumeration of L_{NR} just by taking the second element of each pair in f . Since L_{NR} was stipulated to be non-recursively-enumerable, it follows that f is not computable.

We will use f in the semantics of string terms — the cross-hatch symbol serves as a signal to apply f :

(5) *Semantics of string terms:*

- If α is a number term, then $\llbracket \# \alpha \rrbracket = f(\llbracket \alpha \rrbracket)$

For example, $\#|||$ will denote the third string of L_{NR} in order of increasing length.

Finally, we will let the predicate **TL** hold of any string which is at least two letters in length:

(6) *Semantics of sentences:*

- If α is a string term, then $\llbracket \mathbf{TL}(\alpha) \rrbracket = \mathbf{true}$ if $\text{length}(\llbracket \alpha \rrbracket) \geq 2$; $\llbracket \mathbf{TL}(\alpha) \rrbracket = \mathbf{false}$ if $\text{length}(\llbracket \alpha \rrbracket) < 2$

This completes the syntactic and semantic rules of L .

The semantics just defined is very simple, and fully compositional. I have not bothered here to formalize it algebraically, but an explicit algebraic formalization is given in Lasersohn (2009); it is very easy to define syntactic and semantic algebras for this language, with a homomorphic function mapping every expression onto its denotation.

However, one of the semantic operations for this language is non-computable (namely, the operation deriving denotations of string terms from denotations of number terms). There is no way a person (or a computer) could ever learn that operation — it is non-algorithmic. So there is no possibility that a real person interpreting sentences in this language, or assessing them for truth and falsity, could be applying that operation as part of the interpretation process.

Nonetheless, it is very easy to interpret sentences in this language, and except for a small, finite number of them, it is also easy to assess them for truth and falsity. Consider the following sentence for example:

(7) **TL**(#|||||))

This sentence consists of the predicate **TL**, applying to a string term consisting of the crosshatch symbol followed by 28 vertical strokes. The meaning of this sentence is clear: it means that the 28th string of L_{NR} (when the strings of L_{NR} are arranged in order of increasing length) consists of at least two letters. Moreover, it is clear that this sentence is *true*: I stipulated that L_{NR} was a language over a 26-letter alphabet. So it contains at most 27 strings with length less than 2: The empty string and the 26 single-letter strings. Therefore, the 28th string must consist of at least two letters.

How can we arrive at the truth value, if one of the semantic operations involved in determining that truth value is non-computable? Not by working through a derivation defined directly by the grammar — that is, not by applying the semantic rules to derive denotations for each phrase, based on the denotations of the parts. In fact, it's clear we couldn't be doing that, even apart of considerations of computability: string terms in L denote strings in L_{NR} — but I haven't said exactly which language L_{NR} is. I've given some information about it, but not enough to identify particular strings — in fact, I haven't even said what the alphabet of L_{NR} is; not enough information has been given to determine the denotations of string terms of L even if its semantic operations *were* computable!

There is no great mystery involved. We may not have any way to work through the derivations defined by the grammar, but we do have a lot of background information: We know that L_{NR} has a 26-letter alphabet, we know that f sorts strings by length, and so on. In interpreting sentences of L , and evaluating them for truth and falsity, we can avail ourselves freely of this kind of information and reason about what the denotations of expressions must be like, even if we can't identify them. And that reasoning process is not the same thing as working through a derivation defined by the grammar.

In fact, I would call this *pragmatic* reasoning — “pragmatic” in the sense that it is extragrammatical, uses non-grammatical background knowledge, and forms part of what happens in actual performance, rather than something in the structural description of the language itself. This reasoning process is not in any sense modeled by the grammar; that's not what the grammar is, or what it is for.

Still, this pragmatic reasoning is a process by which we arrive at the propositional content defined by the grammar; it doesn't go beyond that, adding extra information, or anything like that. It's just a different route to the same end. There's no point along that route where the grammar does its thing and then passes the baton to pragmatics; the whole route is pragmatic, but *it is a kind of pragmatics that brings us to a point defined by the semantics*.

That last point is crucial, because the semantics was compositional. And if the pragmatic processes I'm talking about derive the same propositional content as the semantic rules, this

means *they won't produce any counterexamples to compositionality!* The pragmatic reasoning we use to interpret and evaluate sentences in this language may not directly mimic the compositional application of our semantic rules, but the mere fact that we are using pragmatics to construct the interpretation doesn't mean that the language contains any complex expressions whose interpretations aren't determined by the interpretations of their parts.

Of course, this is an artificial example, and a strange one too, because of this appeal to non-computable operations. I am not trying to claim that natural language grammars make use of such operations (though I'm not sure they don't) — that's not the point at all. What I hope is now clear is that just showing that people use pragmatic reasoning and background information to figure out the propositional content of an assertion is not enough to show that the relation of that propositional content to syntactic structure is non-compositional.

Intuitively we think of pragmatics as a top-down, global, non-compositional process; but if we are using pragmatics to figure out what the grammar defines for us, and if the grammar is set up compositionally, then the output of those pragmatic processes will be an interpretation that conforms to compositionality constraints, even if we arrive at it through non-compositional means. In that sense, pragmatics isn't necessarily non-compositional. We can acknowledge that speakers use quite a lot of pragmatics in figuring out the propositional contents of each other's utterances without necessarily viewing that as problematic for compositionality, as long as the grammar is still set up in a compositional fashion.

4. Compositionality and substitutivity. The argument just given is probably not enough to make any anti-compositional contextualists feel worried, because it contains a big "if": *If* the grammar is set up in a compositional fashion, *and* the pragmatic derivation of propositional content is just a way of figuring out what the grammar gives us, *then* the appeal to pragmatics is no threat to compositionality. In my little artificial example, I could just stipulate the grammar, make sure it was compositional, make sure it assigned truth values, so that it determined propositional contents — and the result was that when we reasoned through pragmatically to propositional contents, we arrived at the same ones assigned by the grammar. But in natural language semantics we have to discover the grammar, so we can't just stipulate all those things, and we can't assume in advance that the pragmatic reasoning we engage in during interpretation will just be a different route to a proposition which is also assigned compositionally by the grammar.

What we need at this point is some way of detecting whether the grammar works in a compositional way or not — some way that doesn't involve just pointing out that people have to use pragmatics in order to figure out what the propositional content is, and taking that as decisive. And because the relation between the structure of the grammar and people's actual behavior is not very direct, our argumentation is bound also to be somewhat indirect.

In order to make progress on this sort of issue we need to be a little more explicit about what we mean by "compositionality." As many people have pointed out, different authors take this term as implying different things, and there is a wide variety of different understandings of just what the principle of compositionality is supposed to require. But in most of the more technical or mathematical discussions of the issue, the most common understanding is that "compositional" interpretation means *homomorphic* interpretation.

It does not take a lot of setting up to motivate this basic idea. All we need to assume is that complex expressions are derived from more primitive expressions via some definite set of syntactic operations. Then we say that semantic values are assigned to expressions *compositionally* if there is a function h associating each expression with its semantic value, meeting the following condition:

- (8) For every syntactic operation F , if $F(\alpha_1, \dots, \alpha_n)$ is defined, then there is some operation G such that $h(F(\alpha_1, \dots, \alpha_n)) = G(h(\alpha_1), \dots, h(\alpha_n))$

There are some relatively minor variations on this condition; for example Montague (1970) sets up his system so the outputs of syntactic operations are always defined, rather than only sometimes as I've done here. These kinds of differences won't matter for our purposes as much as the basic idea that what we mean by "compositional" here is *homomorphic*.

That may not be what *everyone* means by "compositional," but it should be recognized that this condition encodes quite directly the idea that the semantic value of a complex expression is determined by the semantic values of its (immediate) parts; we just take the immediate parts of an expression to be the expressions from which it derives via a (single) application of a syntactic operation.

I have deliberately stated the condition here in a way that leaves completely open what we mean by "semantic value" — in particular, whether this is supposed to be denotation, or contribution to propositional content, or linguistic meaning, or whatever — and which leaves open just how all this relates to context. Issues of compositionality arise at all these levels, and the answer may be different at different levels. It is worth bearing in mind that Montague, who really gave the first detailed mathematical development of compositionality as homomorphic interpretation, imposed the homomorphism requirement only at the level of what he called "meaning" — which was closely analogous to Kaplanian character — and not at the level of "sense," which would be where we locate propositional content. But presumably, all these levels will relate to each other in principled ways, so we need to be alert to the possibility that what we say about one of them will constrain what we can say about the others. Our main concern here, I think, is with the level of propositional content, because that is the level that classical contextualist arguments against compositionality seem to be addressing. But we may have to look at other levels as well in order to get a clear view of that one.

In examining the issue of whether semantic values at different levels conform to a principle of homomorphic interpretation, it will be useful to remember a basic mathematical result, which has been known for some time, but which for some reason often seems to get neglected in discussions of compositionality. And that is that:

- (9) Semantic values are assigned homomorphically iff substituting one expression for another with the same semantic value never alters the semantic value of any more complex expression of which it is a part.

I have stated the principle here somewhat informally, but it is an easy exercise to make it more precise and give a real proof. A very detailed mathematical exploration of the relation between

homomorphic interpretation and substitution of identically-valued expressions is given in Hodges (2001).

A corollary of this principle is the following:

- (10) Any counterexample to the claim that semantic values are assigned homomorphically must be presentable in the form of a pair of expressions α , β with the same semantic value, such that substituting α for β in the derivation of some more complex expression γ of which α is a part results in some expression δ with a different semantic value from γ .

Somewhat surprisingly, given how long and how widely it has been known that the principle of homomorphic interpretation is equivalent to the principle of intersubstitutivity of identically-valued parts, semanticists seem quite persistent in offering putative counterexamples to compositionality that aren't of this form (and for which there is no obvious way to recast them in this form) — even when it is quite clear that the notion of compositionality they are working with is essentially just the homomorphism idea, that the semantic value of a complex expression is fixed by the semantic values of its immediate parts.

I will argue that ordinary propositional content is compositional in this sense, so that in principle there are always semantic operations corresponding to our syntactic operations, which derive the contents of complex expressions in contexts from the contents in context of their parts. To do that, I will rely heavily on the intuition underlying this substitutivity principle, and on some basic assumptions about how propositional content relates to truth values.

The first step in arguing that truth *conditions* are determined compositionally is to motivate the intuition that truth *values* are determined compositionally from what is being talked about. Of course what we use an expression to talk about varies from context to context, and so will the truth value of a sentence, so what we really need to motivate is this:

- (11) The truth value of a sentence in a context c is determined compositionally from what the parts of the sentence are used to talk about in c .

The phrase “talk about” is pretty vague, and can be understood in several different ways, not all of which will be relevant here. I will try to sharpen it up as we go, but it will be useful, in getting our intuitions motivated, to put things in this informal, colloquial way.

Since by “compositionally” we basically just mean *homomorphically*, this claim can also be put like this:

- (12) There is a function h associating each sentence with its truth value in c , and each other expression with what it is used to talk about in c , such that for every syntactic operation F , if $F(\alpha_1, \dots, \alpha_n)$ is defined, then there is some operation G such that $h(F(\alpha_1, \dots, \alpha_n)) = G(h(\alpha_1), \dots, h(\alpha_n))$.

But since the principle of homomorphic interpretation is equivalent to the principle that identically valued parts can be substituted without affecting the values of the wholes, this implies the following:

- (13) If two expressions α and β are used to talk about the same thing in context c , and substituting α for β in some sentence φ results in some other sentence ψ , then ψ must have the same truth value in c as φ does.

This is basically just one side of Leibniz' Law, dressed up a little, and I think most of us take it for granted; but at the risk of belaboring the obvious, I would like to think through what the intuition is behind it. I think the basic intuition is the following:

- (14) If you say something true, then swap out part of the sentence and replace it with something else, and the result is that you say something false, then you must not have been talking about the same thing in the part that you swapped out as you were in the part that you replaced it with. (And likewise if you say something false, then swap out part of the sentence and say something true.)

This intuition seems to me to be quite firm, and pretheoretic enough that I don't think it is the byproduct of any particular theoretical perspective, but is really just part of the basic understanding we have as language users. And, of course, this was the fundamental intuition behind Frege's treatment of sentences as denoting truth values, and the intuition behind his claim that in intensional contexts, all expressions denote what would otherwise be their senses. If you accept this intuition, then truth values are determined compositionally from what the parts of the sentence are used to talk about.

I have stated the condition here in terms of "what you are talking about" in a part of a sentence. As I mentioned before, that is a pretty vague way of putting things, and theoretically, we should probably distinguish at least between speaker's reference and semantic reference, and perhaps among other kinds of "talking about" as well. Most semanticists, I think, would claim that the relevant kind of "talking about" here is semantic reference — that is, denotation, or extension. I think that is right, but I also think it is worth noticing that this intuitive justification for compositional determination of truth values is stated in terms of "what you say" and "what you are talking about" — that is, in terms of speech act notions. If we appeal to some notion of semantic reference or extension to try to make this intuition more explicit, we need to remember that it corresponds fairly directly to these speech act concepts, or we will miss some of what the original intuitive explanation was giving us. I am not trying to argue against the speaker's reference / semantic reference distinction; just pointing out that even if we are concerned with semantic reference, it will still have to play some pretty direct role in determination of speech act content.

So far, I have only given arguments for compositional determination of truth *values* in context. But what about truth *conditions*? That's really the level that the main debate is about. But if truth values are determined compositionally, then truth conditions must be too. In particular, given that truth values are determined compositionally from what the parts of the

sentence are used to talk about, truth conditions must be determined from conditions on what the parts of the sentence are used to talk about.

To see this, assume that truth values are determined compositionally from what the parts of sentences are used to talk about. Now, if truth conditions are not determined compositionally from conditions on what the parts of sentences are used to talk about, then there must be two expressions α and β , with identical conditions on what they are used to talk about, but where substituting β for α in some sentence φ results in a sentence ψ with different truth conditions from φ . That is, there must be some condition C under which φ and ψ differ in truth value; but in that same condition C , α and β must be used to talk about the same thing, since they have identical conditions on what they are used to talk about. But then, in C , truth values are not determined compositionally from what the parts of sentences are used to talk about, contradicting our original assumption. So, if truth values are determined compositionally, truth conditions must be as well.

Here we are talking about truth conditions in context — essentially, propositional content, at least if we take propositions to be distinguished only as finely as their truth conditions. There is nothing in what I've just said that stands in the way of the same sentence having very different truth conditions in different contexts, and nothing that requires contextual variation in truth conditions to be tied to the settings of some definite, limited set of parameters, such as might provide the referents of indexicals. But if we accept that truth values are preserved under substitution of coreferring terms, it follows that there operations corresponding to our syntactic operations, and a function mapping expressions onto their contributions to truth conditional content, which is a homomorphism with respect to the syntactic and semantic operations.

This is enough to establish that truth conditional content is determined compositionally — not, of course, by showing that there is anything wrong with contextualism, but rather by showing that compositionality is consistent with contextualism.

5. Invariant grammar with varying operations. My suspicion is that people who are committed to the argument that contextual variation in truth conditions undermines compositionality will not be convinced by the arguments of the last section. The reason is this: Just showing that in any context C there will be a homomorphic mapping from the syntactic algebra to an appropriate semantic algebra does not show that the operations of that semantic algebra are invoked directly by the grammar. After all, we don't have a different grammar for every context; but given everything I've said, we have every reason to expect that the semantic operations used in determining truth conditional content will vary drastically from context to context. So I haven't really shown that truth-conditional content is assigned by the grammar, or that pragmatic contributions to truth conditions are limited to figuring out what the grammar gives us anyway, which is what I suggested earlier.

Could we have a grammar, the same in every context, which assigned truth conditions which vary from context to context, and do it in a compositional way? Compositional assignment of truth conditions in context will require that for each context C there are operations which determine the truth conditions of sentences in C from conditions which hold in C on what the parts of sentences can be used to talk about; but the grammar can't just list these operations, since they will be potentially different in different contexts. But that just means we have to say what

the operations are through some means other than listing. There are all sorts of ways to specify sets of operations, so contextual variation in semantic operations is not automatically problematic for the view that propositional content is assigned compositionally by grammar.

Let me illustrate with an example. Suppose that, for whatever reason, we wanted to build domain-of-quantification effects into the operations which determine the denotations of DPs from the denotations of the determiners and NPs they are composed of. That is, on the analysis in question we would not let determiners or NPs have different denotations in different contexts to reflect the different domains of quantification provided by those contexts; rather, we would build reference to the domain of quantification into the combinatorial rule, something like this:

$$(15) \quad \llbracket \text{Q NP} \rrbracket^C = \llbracket \text{Q} \rrbracket(\llbracket \text{NP} \rrbracket \cap R_C), \text{ where } R_C \text{ is the set of relevant objects in } C.$$

Here, I've deliberately left the C superscript off of $\llbracket \text{Q} \rrbracket$ and $\llbracket \text{NP} \rrbracket$ in order to emphasize that the denotations of the parts of the DP are not considered to vary according to context in this analysis — the contextual variation comes in only at the point where they are combined. (I should probably stress that I am not arguing that this is the correct analysis; I'm just using it to illustrate how a context-invariant grammar can define a different set of semantic operations for each context.)

To make this example concrete, assume that $\llbracket \text{Q} \rrbracket$ is that function F which maps any subset of $\{a, b, c\}$ onto the set of its supersets (so it would be an appropriate denotation for *every* in a model where $\{a, b, c\}$ was the universe). Next, assume that $\llbracket \text{NP} \rrbracket = \{a, b\}$. Now suppose there are two contexts, C_1 and C_2 , with $R_{C_1} = \{a\}$ and $R_{C_2} = \{b\}$. It is easy to see that $\llbracket \text{Q NP} \rrbracket^{C_1} = \{\{a\}, \{a, b\}, \{a, c\}, \{a, b, c\}\}$, while $\llbracket \text{Q NP} \rrbracket^{C_2} = \{\{b\}, \{a, b\}, \{b, c\}, \{a, b, c\}\}$.

Put differently, the rule in (15) defines two different semantic operations O_1 and O_2 , which are used to derive the denotation of Q NP in C_1 and C_2 , respectively. $O_1(F, \{a, b\}) = \{\{a\}, \{a, b\}, \{a, c\}, \{a, b, c\}\}$, but $O_2(F, \{a, b\}) = \{\{b\}, \{a, b\}, \{b, c\}, \{a, b, c\}\}$. But even though the rule defines different operations in different context, the rule itself is context-invariant.

There is no inconsistency, then, between saying that the grammar does not change from context to context and saying that we use different semantic operations in different contexts. The grammar must provide some systematic, context-invariant way of saying what those operations are, but this does not mean that the operations themselves are context invariant.

6. How much pragmatics can a grammar call upon? Advocates of the claim that contextual variation in truth conditions threatens compositionality might respond in at least two ways: 1. They may claim that rules like (15) are not legitimate rules of grammar, since they appeal to pragmatic notions like relevance. 2. They may claim that the contextual effects which threaten compositionality are of a much more thorough-going nature than the effect illustrated in (15), and do not lend themselves to an analogous treatment.

The first objection relies on the idea that grammatical rules may not appeal to pragmatic concepts such as relevance. One might object to such an appeal on the grounds that a grammar

cannot be reasonably expected to say what is relevant in a given context, or to explain the fundamental nature of relevance. All that is the job of pragmatics, not the theory of grammar.

However, it is not necessary, in order to allow rules like (15) in a grammar, for our grammatical theory to include a theory of relevance, or of any other pragmatic notion. Nor is it necessary that the grammatical theory explain how to figure out what is relevant and what isn't (any more than it was necessary for the rules of L in Section 3 to explain how to figure out truth conditions). What *is* necessary is merely that grammatical rules (semantic rules in particular) be allowed to *invoke* pragmatic concepts — relevance, in this case, but conceivably also salience, plausibility, or other notions. These terms would be defined outside the grammatical theory; the grammar wouldn't include instructions on how to figure out when they apply and when they don't, but it could still make use of them.

Perhaps some semanticists would reject the possibility of semantic rules invoking pragmatic concepts because they believe a theory of grammar should be stated entirely in terms internal to the theory itself, and never make use of terms defined outside the theory itself. But that is an implausible condition, which semanticists don't ever conform to in practice: we make use of all sorts of vocabulary for concepts whose theoretical explanation lies outside grammar: terms from set theory, common-sense ontology, etc. Moreover, if we want grammatical theory to have any implications at all for theories of psychology, metaphysics, social interaction, or anything else, (or vice versa) we will need to have some sort of shared vocabulary or translation system between the theories, or we will wind up with a purely abstract, structural theory with no empirical content.

It makes sense, then, that grammatical rules (particularly semantic rules) might make reference to concepts which are explained not as part of the grammar, but as part of a pragmatic theory. If our rules are allowed to invoke such concepts, they can define potentially different semantic operations in different contexts, deriving context-dependent truth conditions in a compositional, homomorphic way.

The second objection was that the contextual effects which threaten compositionality are of a much more thorough-going nature than the effect illustrated in (15), and do not lend themselves to an analogous treatment. The rule in (15) is sensitive to context only because it refers to the set R_C , the set of relevant objects in context C . But R_C seems like a simple contextual parameter; determining this set is not fundamentally different from determining the value of an indexical. But the true force of the contextualist argument against compositionality doesn't come from indexicals, but from the claim that there are contextual effects on truth conditions that *can't* be handled in the same way as indexicals — effects that come from a deep, thorough-going flexibility of use that pervades virtually all expressions, and does not lend itself to explanation in terms of a fixed set of parameters.

However, once we allow grammatical rules to invoke pragmatic notions, there is no need to limit their pragmatic sensitivity to parameter-setting. A rule might appeal to pragmatically-defined concepts in all sorts of ways, up to the point of calling for the application of a semantic operation which is determined entirely by pragmatics, as in the following rule:

$$(16) \quad \left[\begin{array}{c} \alpha \\ \beta \quad \gamma \end{array} \right]^C = \text{the result of combining } [\beta]^C \text{ and } [\gamma]^C \text{ in the way relevant in } C$$

Personally, I doubt whether natural languages could ever use rules like this. If an entire language worked this way, we could almost say that it had no semantics at all, but only syntax and pragmatics!

We could *almost* say that, but not quite. The rule given in (16) is still a semantic rule, and does assign denotations to trees, relative to contexts. It doesn't tell us how to *figure out* what that denotation is — but that is not what semantic rules should be expected to do anyway. Moreover, it assigns denotations in a completely compositional, homomorphic way, based on syntax. Allowing rules like this may be threatening to the idea that formal, grammar-based semantics bear most of the *explanatory burden* in the theory of truth conditions, but it is not threatening to the idea of compositional interpretation, or to the idea that in principle, contents are assigned as part of grammar.

What *would* be threatening to compositional assignment of truth conditions? Only a case where substituting one expression for another, used in the same context to talk about the same thing, resulted in a different truth value for the sentence as a whole. We would allow such an effect if we had rules like the following, for example:

$$(17) \quad \left[\begin{array}{c} \alpha \\ \beta \quad \gamma \end{array} \right]^C = \text{the thing which is relevant to } \alpha \text{ in } C$$

This rule frees up the content of α from any requirement that it be related to the contents of its constituents in any way (unless “relevant to α ” somehow implies a connection to β and γ). Interpretation is done purely by pragmatics, with no contribution from syntactic structure. But even the most ardent anti-compositionalists don't go this far.

7. *Contextual variation actually saves compositionality.* Is there anything short of the unrealistic extreme illustrated in (17) that might counterexemplify compositionality? We might reason as follows: A homomorphism has to be a function, and a function can assign only one value to each argument, so homomorphic interpretation seems to preclude ambiguity. Of course ambiguity is rampant in natural language, so a principle of homomorphic interpretation would seem to be problematic.

However, if we are concerned with content *in context*, there is much less ambiguity than if we were concerned with ambiguity out of context — especially if by ambiguity we don't mean cases where it is difficult to tell which of several possible meanings is the “correct” one in context, but rather, cases where more than one meaning is “correct.” Probably the main examples would be puns and deliberate double-entendres. But in these cases it seems that an expression is used to

talk about more than one type of thing at once; for example in *Police were called to a daycare center where a three-year-old was resisting a rest*, the phrase *a rest* is used both to talk about repose, and about being stopped by police. Substituting what would ordinarily be a synonym (such as *repose*) and noting that the substitution does not preserve meaning would therefore *not* constitute an authentic counterexample to the claim that truth values are determined compositionally from what the parts of a sentence are used to talk about — hence not a counterexample to the claim that truth *conditions* are determined compositionally from *conditions* on what the parts of the sentence are used to talk about.

Another class of apparent counterexamples involve cases where the truth conditions of a sentence depend on phonological or other metalinguistic features of one of the constituents: *Giorgione is so-called because of his size*, etc. But in these cases, part of what we are talking about is the language itself, so substitution of what would be a synonym in other contexts should not be expected to preserve truth values, even if compositionality is assumed.

Likewise, the classic Fregean examples involving intensional contexts can be dealt with just as Frege did, by claiming that words in these contexts aren't used to talk about what they are used to talk about in other contexts. Despite Davidsonian claims that “innocent” speakers don't suspect such a reference shift, I find that it remarkably easy to convince inexperienced undergraduates that sentences like *The ancient Babylonians believed that the morning star is the evening star* is not about the planet Venus, but about the ancient Babylonians' conception of the planet Venus — that is, about a mode of cognitive presentation.

In all these cases — puns, double entendres, metalinguistic interpretations, and intensional contexts — we find that substitution of what would ordinarily be a synonym fails to preserve truth value, precisely *because* in these contexts words that are ordinarily used to talk about the same thing are instead used to talk about different things. It is this contextual shift in interpretation which allows us to preserve the idea of homomorphic interpretation. In other words, far from being problematic for compositionality, contextual variation in interpretation is precisely what *rescues* the claim that interpretation is compositionally assigned from apparent counterexamples.

8. Conclusions. The following may serve as a summary of the major points advocated in this paper:

- We must carefully separate the issue of how propositional content relates to syntactic structure from the issue of how people figure out what proposition a sentence expresses.
- The fact that people employ pragmatic reasoning in the latter task does not imply that the relation of content to structure is non-compositional.
- Compositionality (in the sense of homomorphic interpretation) is equivalent to a familiar substitutivity principle.
- This principle is motivated by the intuition that substitutions which affect truth value must be of expressions which are being used to talk about different things.

- Arguments against compositionality must address this intuition, justify setting it aside, and explain why we have it.
- Compositional determination of truth values implies compositional determination of truth conditions.
- Semantic rules which form part of grammar may reasonably be expected to invoke pragmatic concepts like relevance, salience, etc.
- This does not mean that grammatical theory is responsible for defining or explaining these concepts, nor does it represent a relabeling of pragmatic theory as part of grammar.
- Compositional assignment of contents by a grammar is compatible with radical pragmatic effects on truth conditions.
- Even if the only way to identify the propositional content of a sentence in context requires heavy appeal to pragmatics, this does not show that this content violates compositionality principles.
- Contextual variation in truth conditions actually *saves* the compositionality thesis from apparent counterexamples.
- *Pragmatics can be compositional!*

References

Carston, Robyn (1988) 'Implicature, explicature, and truth-theoretic semantics', in *Mental representations: The interface between language and reality*, 155-181, ed. by Ruth M. Kempson, Cambridge University Press.

Hodges, Wilfrid (2001) 'Formal features of compositionality', *Journal of Logic, Language and Information* 10.7-28.

Lasnik, Howard (2009) 'Compositional interpretation in which the meanings of complex expressions are not computable from the meanings of their parts', in *Theory and evidence in semantics*, 133-158, ed. by Erhard Hinrichs and John Nerbonne, CSLI Publications.

Montague, Richard (1970) 'Universal grammar', *Theoria* 36.373-398.

Recanati, François (2005) 'Literalism and contextualism: some varieties', in *Contextualism in philosophy: knowledge, meaning and truth*, 171-196, ed. by Gerhard Preyer and Georg Peter, Oxford University Press.

Travis, Charles (1997) 'Pragmatics', in *A companion to the philosophy of language*, 87-107, ed. by Bob Hale and Crispin Wright, Blackwell Publishers.