

THE DETERMINATION OF CONTENT

Zoltán Gendler Szabó
Yale University

Abstract: I identify a notion of compositionality at the intersection of the different notions philosophers, linguists, and psychologists are concerned with. The notion is compositionality of expression content: the idea that the content of a complex expression in a context of its utterance is determined by its syntactic structure and the contents of its constituents in the contexts of their respective utterances. Traditional arguments from productivity and systematicity cannot establish that the contents of linguistic expressions are compositionally determined in this sense. I present a novel argument for this thesis based on plausible premises about literal use and a detailed defense of the compositionality of speech-act content.

Jerry Fodor writes: “So non-negotiable is compositionality that I’m not even going to tell you what it is.”¹ He has a point. Some sort of compositionality is beyond dispute: there must be a recursive mechanism underlying linguistic competence. At the same time, it has proven exceedingly difficult to state the principle precisely without falling into triviality or falsehood. Despite this, I think we should keep trying.

The need for clarity is underlined by the fact that in the paper which contains the quote above, Fodor actually argues *against* the compositionality of natural language. He does so in order to argue *for* the compositionality of human thought. He reasons as follows: either thought or language must be compositional, and if language isn’t thought must be. This, in turn is supposed to show that thought is explanatorily *prior* to language: for whichever of the two is compositional must be what is meaningful in the first place.

Fodor’s argument that natural languages are not compositional is based on the alleged empirical fact that “sentences are remarkably inexplicit with respect to how the thoughts they express are put together.”² Fodor thinks the claim is untendentious; he illustrates it with only two examples. The first is an utterance of ‘it is three o’clock’ made at 3pm which, he says, is inexplicit as a way of expressing the thought that it is three o’clock *in the afternoon*. The second is an utterance of ‘the book is on the table’ made in the presence of multiple tables with books which, he says, is inexplicit as a way of

¹ Fodor (2001): 6.

expressing a singular thought about a particular book on a particular table. As Fodor is well aware, these are not examples that tend to sway semanticists. But he thinks this is only because “it’s the house rule in formal semantics that you hold onto the compositionality of natural language, whatever the cost in face implausibility may be.”³ Fighting words, but they do contain a kernel of truth. Working semanticists often claim that compositionality is a methodological principle, but they hardly ever assess the costs and benefits of adhering to it. Compositionality has the whiff of a dogma, not of a maxim.

There is a term in the literature for the considerations Fodor presents against the compositionality of natural language: they are called *underdetermination arguments*.⁴ Underdetermination arguments aim to show that a sentence with a fixed content can be used to assert multiple contents. Regarding the two examples Fodor mentions, a proponent of underdetermination would say that in uttering ‘it’s three o’clock’ or ‘the book is on the table’ the speaker can actually *assert* (not only *convey* by means of an implicature) the proposition that it’s three o’clock in the afternoon and the proposition that *b* (some particular book) is on *t* (some particular table), respectively. These propositions are asserted, even though they are not the compositionally determined semantic values of the sentences uttered. Such a proposal saves the compositionality of sentence content by sacrificing the compositionality of assertion content.⁵ In this paper, I argue that this strategy is mistaken.

Here is the plan. I begin by spelling out more precisely what the principle of compositionality of meaning is and then, in section 2, what the principle compositionality of sentence content (and more generally, *expression content*) amounts to. I argue that traditional arguments for compositionality – productivity and systematicity – cannot

² Fodor (2001): 12.

³ Fodor (2001): 13.

⁴ Such arguments have gained a lot of momentum in the last two decades. Here are a few milestones in the spread of the idea: Searle (1978), Travis (1985), Sperber and Wilson (1986), Perry (1986), Recanati (1989), Bach (1994), Levinson (2000), Carston (2002), Cappelen and Lepore (2005), Soames (2005).

⁵ So, unlike Fodor, proponents of underdetermination arguments tend to think that these examples do not show that *natural language* is non-compositional. Important exceptions are Lahav (1989) and Travis (1994). The latter will be discussed in detail in section 4 below.

establish that the latter principle holds. In section 3, I turn to compositionality of assertion content (and more generally, *speech-act content*), and I show that if speech-act content is compositional, so is expression content. In section 4, I show how to respond to underdetermination arguments against the compositionality of speech-act content, and argue that the responses constitute strong empirical support for speech-act compositionality.

My argument is bound to strike some as odd: it proceeds from a widely rejected premise (that speech-act content is compositional) to a widely accepted conclusion (that expression content is). But if I am right both the rejection and the acceptance are based on false assumptions.

1. Compositionality

The principle of compositionality, as I propose to use this term, is the claim that the meaning of a complex expression depends on two things only: how its constituents are put together and what those constituents mean.

- (C) The meaning of a complex expression is determined by its structure and the meanings of its constituents.

This is bottom-up determination claim. It is formally analogous to the claim that the physical properties of an ordinary middle-sized dry object are determined by the way it is composed from elementary particles and the physical properties of those particles.

(C) is a thesis about some particular language. Questions of constituency and structure are settled by the syntax of the language in question; the meanings of simple expressions are assigned by its lexicon. Facts of syntax and lexicon must be respected; otherwise we run the danger of trivializing compositionality. As long as we are talking about *English*, we cannot identify ‘John kissed’ as a constituent of ‘John kissed Mary’, for standard syntactic tests reveal that it is not. Similarly, we cannot claim that ‘green’ means

something different in ‘green apple’ and in ‘green room’, for there is no evidence from lexical semantics that can back up existence of such an ambiguity.⁶

Compositionality is often seen as a claim that says that there is a *function* from structure and the meanings of parts to the meanings of wholes. But functions are cheap and determination is not – there is probably a function from the GDP of each country to the number of its bald citizens but the former surely does not determine the latter. Obviously, (C) says *more* than that there is such a function. However, since it is a matter of controversy *how much more* it says, for the purposes of this paper I won’t assume any strong construal of compositional determination.⁷

There is a key ambiguity in (C) that has long been neglected in the literature: the plural definite description in it can be construed either distributively or collectively. The collective reading is compatible with the possibility that the meaning of certain complex expressions depend not only on the meanings their constituents have in themselves, but also on relations that hold among these meanings. I will follow the usual practice in reading ‘the meanings of its constituents’ in (C) distributively.⁸ Thus construed (C) (together with certain straightforward assumptions about the syntax of the language under consideration⁹) entails that ambiguity in a complex expression is either syntactic (derives from the fact that the expression has more than one structure) or lexical (derives from the fact that some lexical constituent of the expression has more than one meaning).

⁶ Compositionality is a demonstrably empty principle if one neglects syntactic and lexical constraints. Janssen (1983) has a proof that we can turn any meaning assignment on a recursively enumerable set of expressions into a compositional one, as long as we can replace the syntactic operations with different ones. Zadrozny (1994) has shown that this can also be done by replacing the old meanings with new ones from which they are uniformly recoverable.

⁷ In Szabó (2000a), I argue that (C) is best construed as the claim that there is a single function *across all possible human languages* from the structure of any complex expression and the meanings of its constituents to the meaning of that complex expression.

⁸ Fine (2007): 26 distinguishes between “compositionality proper” and “intrinsicism”. The former is the collective reading of (C); the latter is the claim that “the collective meaning of the component expressions is exhausted by their individual meanings.” As Fine notes, the standard construal of (C) is the distributive one, which is the result of combining the collective reading with the intrinsicist thesis.

⁹ The usual assumption is that syntax defines a *syntactic algebra*; cf. Janssen (1983).

(C) is not a principle anyone is interested in for its own sake. Philosophers, linguists, and psychologists are all concerned with substantially stronger theses.

The *philosopher's* compositionality is a claim about explanatory *priority*. It says that complex expressions are meaningful because their parts are:

(Φ) Complex expressions have their meanings in virtue of their structure and the meanings of their constituents.

According to Frege's *context principle*, words have their meanings in virtue of the meaning of some sentence (or sentences) in which they occur as constituents. Those sentences then would presumably not have their meanings in virtue of their structure and the meanings of their constituents – otherwise we would have circular explanations of meaning. So, the context principle is incompatible with (Φ). But it is not incompatible with (C): unlike priority, determination can be symmetric.^{10,11} (According to Newton's second law of motion, for example, given the mass of an object, the applied force and the acceleration mutually determine each other.)

The *linguist's* compositionality requires a tight connection between syntax and semantics:

(Λ) The meaning of a complex expression is determined by its immediate structure and the meanings of its immediate constituents.

(C) does not speak of immediate structure and constituents: all it says that the meaning of a complex expression is determined by its *entire* structure and the meanings of *all* its constituents. In assigning meaning to a complex expression (C) allows us to look at the

¹⁰ There is a long-standing debate whether Frege's commitment to the context-principle can be reconciled with his commitment to compositionality; cf. Janssen (2001) and Pelletier (2001) for recent opinions on the matter. As long as the former is understood as a claim about priority and the latter as a claim about determination, I see no conflict.

¹¹ The fact that Fodor argues for (Φ), not merely (C) is clear at many places where he discusses compositionality. (He sometimes states the principle as saying that the meaning of a complex expression is *inherited* from the meanings of its parts; cf. Fodor (1991): 14.) He also needs (Φ), not (C) in his argument (mentioned in the second paragraph of this paper) that thought is explanatorily prior to language.

meanings of constituents deep down the tree representing its syntactic structure, while (Λ) permits only looking down one level.¹²

The *psychologist's* compositionality is a thesis about the way we understand complex expressions:

(Ψ) We understand a complex expression by understanding its structure and its constituents.

(C) guarantees that one could – setting aside cognitive limitations – determine the meanings of complex expressions bottom-up. According to (Ψ), this is in fact how we proceed. The fact that we understand complex expressions we never heard before is a reason to think that *sometimes* we do understand complex expressions in this way, but it does not follow that we *always* do so. Moreover, it might even be the case that we understand certain complex expressions *regularly* by tracing a structure different from their syntactic structure.¹³

What makes (C) an interesting thesis is that it is a common consequence of (Φ), (Λ), and (Ψ). If it is false, all these important claims are false too. But to assess the truth-value of (C) we need to be more explicit about what we mean by ‘meaning’.

2. Expression content

Meaning is what speakers grasp when they understand an expression. There are at least two notions of understanding, and correspondingly, at least two notions of meaning.

¹² The semantics for attitude ascriptions in Carnap (1947) satisfies (C) but not (Λ). It says that any two tautologies have the same semantic value but the semantic values of sentences ascribing tautological beliefs to the same subject can differ. This is due to the fact that tautologies may differ in their structure, and consequently, in the semantic values of their constituents as well.

¹³ For examples of languages for which (C) holds but (Ψ) fails see Szabó (2000c): 77 – 80. Horwich (1998): 155 thinks that something like (Ψ) is true by definition. This strikes me as a rather unfortunate stipulation in that it seems neither to accord with the ordinary meaning of ‘understanding’, nor to create a theoretically useful novel term.

There is a clear sense in which any competent speaker of English understands the sentence ‘she went inside’ and there is an equally clear sense in which only those do who know who went where. Following David Kaplan, we can say that those who understand the sentence in the first sense know its *character*, and those who understand it in the second know its *content*.¹⁴

Kaplan advocates a specific conception of character and content. He sees content as an object, a property, a relation, or a complex built up from such building blocks. He takes context to be an ordered triple of a speaker, a time and a place. And he defines character as a function from the context determined by a particular utterance of an expression to its content as uttered. I do not wish to commit myself to any of these specific conceptions – doing so would needlessly reduce the generality of my discussion. All I mean by character is *linguistic meaning out of context* and all I mean by content is *linguistic meaning relative to context*. I also wish to be fairly liberal about what belongs to the context relative to which content is determined. Following David Lewis,¹⁵ I acknowledge that an adequate semantics for a natural language *may* have to include in the context features beyond the components of Kaplan-contexts: an addressee (to interpret ‘you’), objects available for demonstration (to interpret ‘this’ and ‘that’), a segment of discourse that has taken place (to interpret anaphoric expressions), orientations (to interpret ‘left’ or ‘down’), standards of precision (to interpret ‘hexagonal’ and ‘flat’), and much else. In fact, given the open-endedness of the list of factors that may influence the assignment of content, it is probably best not to think of contexts as ordered *n*-tuples at all. My favored view is that the context of an utterance is part of the common ground between the speaker and the addressee, but this won’t be assumed either in what follows.¹⁶

The idea that semantics must be *two-tiered* because it needs to assign both character and content is widely but not universally accepted. It is sometimes argued that semantics

¹⁴ Cf. Kaplan (1977).

¹⁵ Lewis (1970).

¹⁶ According to Stalnaker (2002), a proposition *p* is part of the common ground in a conversation iff the speaker and the addressee accept *p*, they both believe that they accept *p*, they both believe that they both believe that they accept *p*, ... and so on. Context should not be thought of as *all* of the common ground,

should restrict itself to the assignment of character. So, for example, Jim Higginbotham recommends that the semantic clause for ‘she is lazy’ should be following conditional: ‘if x is referred to by ‘she’ in the course of an utterance of ‘she is lazy’ and x is female, then that utterance is true iff x is lazy.’ In this way, he hopes to stay clear of the “morass of communicative context.”¹⁷ One way to do this systematically would be to employ in our semantic theory conditionalized schemata whose antecedents specify an arbitrary assignment to all the context-dependent lexical items in the language. This way semantics reveals only the character of a linguistic expression – its content can be determined by ascribing contents to the context-dependent lexical items and by applying modus ponens with the relevant semantic clause as the main premise.

The near-opposite view has also been advocated. Jeff King and Jason Stanley think ascription of character to complex expressions is otiose. Given the standard notion of character (function from contexts to contents), “both a semantics that assigns characters to simple expressions and recursively assigns characters to complex expressions *and* a semantics that assigns characters to only simple expressions allow for an assignment of the same contents in contexts to simple and complex expressions.”¹⁸ And since assigning character to simple expressions only is arguably simpler than assigning it to both simple and complex expressions, perhaps we should choose the former option.

Both of these proposals make a crucial assumption. They take it for granted that the context-dependence of a complex expression ultimately derives from the context-dependence of its lexical items. Only lexical items are mentioned in the antecedents of Higginbotham’s semantic clauses and only lexical items receive a context-sensitive semantic value for King and Stanley. This means that semantic theories that acknowledge context-dependence that emerges above the lexical level are automatically ruled out. While it is true that such theories violate a version of the compositionality principle, their summary dismissal is unjustified. Let me elaborate.

since then context would include linguistic meanings as well. Such a broad conception of context would make all expressions trivially context-dependent.

¹⁷ Higginbotham (1988): 40.

¹⁸ King and Stanley (2005): 128.

The principle of compositionality of character is easy to state; all we need to do is to replace occurrences of the word ‘meaning’ in (C) with occurrences of ‘character’. As a first approximation, the principle of compositionality of expression content can be stated analogously, if we relativize to context: the content of a complex expression relative to a context of utterance is determined by its structure and the contents of its constituents relative to that context. But this is too simple – it assumes that the context of utterance is the same for all constituents of the sentence, which in general need not be the case. I can truthfully utter ‘it is hot here but not here’ as long as I am walking from one room to another; in this case the two occurrences of ‘here’ are uttered in different contexts. Taking these complications into account, the principle can be stated as follows:

(C_{ec}) The content of a complex expression relative to a context of its utterance is determined by its structure and the contents of its constituents relative to the contexts of their respective utterances.

(C_{ec}) entails that any context-dependent complex expression must have at least one context-dependent constituent and that all context-dependence derives from lexical context-dependence.¹⁹ This is what Higginbotham as well as King and Stanley presuppose. An example will help make clear what sorts of theories are ruled out by (C_{ec}).

According a suggestion due to Jeff Pelletier, the content of ‘every man’ varies from context to context depending on what is included in the domain of quantification. At the same time, Pelletier maintains that the contents of ‘every’ and ‘man’ are fixed.²⁰ This is in clear conflict with (C_{ec}). Pelletier claims that compositionality, properly construed, requires only that the content of a complex expression be determined by the contents of its constituents and the way those contents – not the expressions themselves! – are combined. He assumes that the manner in which the contents of ‘every’ and ‘man’

¹⁹ Cf. Stanley (2000). In Szabó (2000b) I called this the *context principle*.

²⁰ Pelletier (2003).

combine varies with the context, and this is how he seeks to account for the contextual variation of the domain of quantification associated with ‘every man’.²¹

I don’t accept Pelletier’s proposal, but I also don’t think it can be dismissed without further ado. Commitment to the compositionality of character is one thing; commitment to the compositionality of expression content is another. Traditional arguments from the productivity and systematicity of understanding give us some grounds to accept the former, but no grounds for accepting the latter.

The argument of productivity goes as follows. We start from the observation that competent speakers can understand complex expressions they have never encountered before. Since this is so, there must be something competent speakers know (perhaps tacitly) on the basis of which they can determine what those complex expressions mean. What can this knowledge be? It seems that the only thing we can plausibly point at is knowledge of the structure of complex expressions and knowledge of the meanings of their simple constituents.²²

This sort of argument does not even get off the ground when the relevant notion of meaning is that of expression content. It is simply not true that competent speakers can in general understand – know the content of – complex expressions they never encountered before. In fact, they don’t know in advance the content of familiar complex expressions either, as long as they are context-dependent. Linguistic competence alone is not enough to determine the content of ‘she went inside’.

The argument from systematicity starts with a more contentious observation: that anyone who understands some complex expressions understands all other complex expressions

²¹ Pelletier’s proposal has viable alternatives that do not violate (C_{cc}). Perhaps ‘every man’ always expresses unrestricted quantification over all men, perhaps the content of ‘man’ in a context includes only the men who are contextually salient, perhaps the content of ‘every’ includes a domain of quantification. I favor the second alternative; cf. Stanley and Szabó (2000).

²² What sort of knowledge (if any) is involved in one’s understanding a word or syntactic structure is a matter of considerable disagreement. Wittgensteinians tend to argue that it is a certain practical ability or knowledge-how, Davidsonians tend to see it as a theoretical disposition or knowledge-that. Intermediate

that can be built up from their constituents using syntactic rules employed in building their structures. (To use Jerry Fodor's favorite example, anyone who understands 'brown cow' and 'black cat' also understands 'black cow' and 'brown cat'.) Since this is so, there must be something competent speakers know (perhaps tacitly) on the basis of which they can determine what the complex expressions built through recombination mean. What can that knowledge be? The only plausible assumption seems to be that it is knowledge of the structure of the original complex expressions and knowledge of the meanings of their simple constituents.

This argument looks more promising when it comes to expression content. Someone who knows the content of 'he is happy' and the content of 'she is sad' relative to a context must know who 'he' and 'she' refer to in that context and will presumably also know the content of 'she is happy' and 'he is sad'. But when we consider more complex cases, such as cases relevant for the assessment of Pelletier's proposal, the situation becomes less clear. Someone who knows the content of 'every man is happy' relative to a certain context of utterance must know which men belong to the domain of quantification: they are the ones who must be happy if the sentence is true. Is this, plus knowledge of the content of 'Sally is a woman' relative to the same context enough to guarantee that this person also knows which women belong to the domain of 'every woman is happy' relative to that context? It is hard to tell – I certainly have no clear intuitions about the matter.

I conclude that to the extent that the traditional arguments give us reason to believe in compositionality, they support the claim that the character of a complex expression is compositionally determined, not the claim that its content is. What is productive and systematic is our grasp of meaning out of context. To know the content of a linguistic expression requires more than linguistic competence – we must also be able to identify the relevant features of the context. The difference between character and content does not matter as long as we are focusing on semantic theories for a fragment of language that

positions are also possible – e.g. Paul Horwich thinks knowledge of word meaning is an *implicit* knowledge-that constituted by use. I take no stand on this issue here.

contains context-independent expressions only, or if we knowingly abstract away from context-sensitivity in semantic theorizing. But it does matter when context is at issue.

One moral is that in our most general characterization of the task of semantic theories we should not presuppose (C_{cc}), as Higginbotham or King and Stanley do. The other is that those of us who believe in this principle – apparently the vast majority of people working on natural language semantics – need a new argument.

3. Speech-act content

When a sentence is uttered with communicative intent, the utterance itself – a particular event – has certain content. This content may or may not be the same as the content of the sentence relative to the context of utterance. Suppose A were to address B with the following: “You are my sunshine, my only sunshine, you make me happy when skies are grey.” The sentence is among other things about sunshine and the color of the skies, but – assuming that the context is ordinary and that A is earnest – the utterance is not. All *the speaker* said was something about his feelings towards the addressee. Given that he was not speaking literally, what he said can be true even if the addressee bears no resemblance to electromagnetic radiation and even if the speaker has a splitting headache whenever it rains.

The content of one’s utterance of a sentence is the content of the speech act one performs in uttering the sentence.²³ In uttering sentences we typically perform *illocutionary acts* – we assert, command, question, request, we make a bid, a promise, a guess, etc.

Illocutionary acts are moves in a conversation. Making an illocutionary act involves – using Austin’s original phrase – “conventional consequences” such as, assuming and

²³ One may well perform multiple speech acts in uttering a single sentence. Nonetheless, I assume that in all but a negligible percentage of those cases (e.g. the case of the famous utterance of ‘Peccavi’ discussed by Paul Grice) there is one *principal* speech-act performed. For example, in uttering ‘Fortunately, Bill did not come’ one arguably asserts both that Bill did not come and that this is fortunate. The former is the principal speech-act, which can be seen from the fact that the response ‘That’s not true!’ is naturally taken to express that Bill did not come. When I speak of speech-act content, in cases when multiple speech-acts are performed by uttering a single sentence, I mean the content of the principal speech-act.

conferring rights and obligations. We can perhaps perform illocutionary acts in uttering individual words or phrases (e.g. when we hold up a letter and utter ‘From Spain’²⁴) but uttering expressions *within* larger expressions does not typically constitute an illocutionary act. If I assert a disjunction I do not thereby assert the disjuncts.

Still, uttering an expression – *any* expression – with communicative intent is a speech act of some sort, and a speech act – *any* speech act – has content of some sort. To deny that there is a semantic difference between the utterance of the letter ‘s’ and the utterance of the word ‘snow’ within an utterance of the sentence ‘snow is white’ is to deny the obvious. There are distinctively semantic questions that can be raised about the utterance of ‘snow’ – for example, we can ask what the speaker of the utterance was referring to in making this utterance – but no distinctively semantic questions that can be raised about the utterance of ‘s’. It is unclear to me how one could account for this contrast short of allowing that, unlike the utterance of ‘s’, the utterance of ‘snow’ is a genuine content-bearing speech-act.²⁵ If I assert a disjunction, I do perform a non-illocutionary speech act in uttering each disjuncts.

How is speech-act content determined? A straightforward idea would be that the determination is compositional:

(C_{sc}) The content of the speech-act performed by the speaker in uttering a complex expression is determined by its structure and the contents of the speech-acts performed by the speaker in uttering its constituents.

Clearly, the content of the speech-act performed by a speaker uttering a complex expression depends on these two factors. The null assumption should be that it does not depend on anything else besides. In case of “You are my sunshine, my only sunshine, you make me happy when skies are grey” (C_{sc}) would tell us that whenever these lines are used to make an utterance whose content differs from the content of the sentence, it must

²⁴ See Stainton (2006) for an extended argument that sub-sentential expressions can be used to make assertions.

²⁵ In acknowledging non-illocutionary speech-acts I follow Searle (1969), which already identified referring and predicating as such. Searle calls these *propositional speech-acts*.

be that some of the contents of some of the utterances of constituent words (including, presumably, ‘sunshine’, ‘skies’ and ‘grey’) also differ from the contents of those words. This, I think, conforms reasonably well to ordinary intuitions.

One might wonder what happens when more elaborate or less common metaphors are used by a speaker. Instead of sunshine, one might call one’s beloved a bag of myrrh and instead of invoking the grayness of sky, one might characterize one’s discontent by invoking a desolate garden. Are these not counterexamples to the compositionality of speech-act content? I don’t think so. These are cases where the content of the speech-act is what it seems – that one’s beloved is a bag of myrrh or that one’s mood is a desolate garden. What is unusual is the *force* of these utterances: they are not assertions but pretend assertions – speech-acts when a speaker acts as if he asserted some content.²⁶ If there are counterexamples to (C_{sc}), they must be of a fundamentally different sort. Here is a representative example from Charles Travis:²⁷

As an arbitrary example, consider the words ‘The leaf is green’, speaking of a given leaf, and its condition at a given time, used so as to mean what they do mean in English. How many distinct things might be said in words with all that true of them? Many. That emerges when we note that one might speak either truth or falsity in such words, if the leaf is the right way. Suppose a Japanese maple leaf, turned brown, was painted green for a decoration. In sorting leaves by colour, one might truly call this one green. In describing leaves to help identify their species, it might, for all the paint, be false to call it that. So words may have all the stipulated features while saying something true, but also while saying something false.

Travis makes two claims: that the sentence ‘the leaf is green’ can be used to say different things, and that the source of this ambiguity is not lexical. (He also presupposes, plausibly, that we have no syntactic ambiguity here.) The words are “used to mean what they do mean in English” – that is, the contents of their utterances are the same in the two contexts described. If Travis is right about this, the contents of the speech-acts the speaker performs in uttering certain words underdetermine the content of the speech act she performs in uttering a sentence those words make up.

²⁶ Cf. Walton (1990).

²⁷ Travis (1994): 171-2.

Travis says this is an arbitrary example and he is right in suggesting that many others can be constructed following the same recipe. Consider the sentences ‘Steel is not strong enough’, ‘That flea is small’, and ‘All students passed’. Each of these can be used to assert different things in different contexts – e.g. that steel is not strong enough to support the roof or that it is not strong enough to crush a walnut, that the flea in question is small for a flea or that it is small for an animal, and that all students in one teacher’s class passed or that all students in another’s class passed. In each of these cases, one can appeal to intuition to argue that the utterances of the words share their contents in the contexts under consideration, and consequently, that all three yield counterexamples to (C_{sc}).

The standard view is that this style of argument is persuasive – speech-act content is non-compositional, even if expression content is. I think the standard view is mistaken. There is reason to think that in all these cases we have a shift in the content of an utterance of a constituent word corresponding to the shift in the content of the utterance of the sentence. In the next section I present an argument to this effect.

Suppose for a moment that my argument in the next section succeeds. We would then not only be in a position to defend (C_{sc}), we could also give an argument for (C_{ec}). The argument rests on two additional premises:

- (L₁) When an expression is used literally all its constituents are.
- (L₂) Any expression can be used literally in any context.

(A linguistic expression is used literally just in case the content of its utterance is identical to its own content in the context of utterance.) The argument goes as follows. Consider an arbitrary complex expression e and an arbitrary context c . According to (L₂), e can be used literally in c . Suppose it is – then the content of e in c is identical to the content of the speech-act of e ’s utterance. According to (C_{sc}), this content is determined by the structure of e and by the contents of the speech-acts performed uttering e ’s constituents. According to (L₁), the constituents were also used literally, so the contents of the speech-acts performed uttering the constituents are identical to the contents of the constituents within their respective contexts of utterance. So, the content of an arbitrary

complex expression in an arbitrary context of utterance is determined by its structure and the contents of its constituents in the contexts of their respective utterances. In other words, (C_{ec}) holds.

(L₁), I take it, is beyond dispute. Non-literal uses cannot neutralize each other: if I use a word in a sentence non-literally, there is no complementary non-literal use of another word in the same sentence that will result in a literal use of the sentence as a whole. (L₂) is also hard to reject. What could justify the claim that the content of a linguistic expression in a certain context is such-and-such, if no utterance of that expression in that context can have that content?^{28,29}

The success of this argument depends on the plausibility of (C_{sc}), which in turn depends on whether Travis's challenge and its ilk can be properly responded to.

4. An argument for the compositionality of speech-act content

Consider the two situations Travis discusses. In the first people are sorting leaves by color for the purpose of decoration, in the second they are trying to identify what species various leaves belong to. In the first situation, A holds up the red maple leaf painted green and utters the sentence 'the leaf is green' and thereby says something true. In the second, B holds up the same leaf, utters the same sentence, and thereby says something false. We are invited to conclude that they said different things, and I will go along with this suggestion.³⁰ What is the source of the difference? Travis says it cannot be linked to

²⁸ Note that (L₂) is compatible with the existence of expressions that are often – or normally, or always – used non-literally, as long as they *can* be used literally. (Bach (1994) argues that the standard use of some expressions is non-literal.) It is also compatible with the view that certain sentences express a propositional matrix – something that presumably cannot be the content of illocutionary speech-acts – as long as they can be used literally in non-illocutionary speech-acts. (Soames (2005) argues that sentences containing proper names express propositional matrices.)

²⁹ (L₂) is trivial if one assumes (i) that speaker-intentions are independent of context and (ii) that speaker-intentions determine whether one's use of a linguistic expression is literal. But I think (L₂) should be attractive even for those who reject (i) or (ii).

³⁰ One may reject this conclusion. Perhaps A and B said the very same thing, something that is true in the context of A's utterance, but false in the context of B's utterance. (This would be the sort of view John MacFarlane has called *non-indexical contextualism* applied to speech-act content; cf. MacFarlane (forthcoming).) But if this is so, Travis's argument against (C_{sc}) doesn't even get off the ground.

the lexicon – utterances of all the words carry their usual fixed contents. I now present an argument that shows that this is not true.

Imagine that someone who did not hear A's utterance well asks the question 'what did you say this leaf was?' Suppose A answers by uttering the sentence 'it is green'. Since A's response is a *restatement* it seems clear that in calling the leaf 'green' A said the same thing he did in the first time.³¹ In other words, the contents of his two utterances of the word 'green' are identical. And of course, if B answers the question 'what did you say this leaf was?' by uttering 'it is green', the contents of her two utterances of the word 'green' are also the same. Consider now the utterances of 'it is green' by A and B. Both are about the painted leaf but the first is true and the second false, so in uttering this sentence A and B said different things about the same leaf. But whatever was said about the leaf in uttering 'it is green' was said in calling it 'green'. So, the contents of the utterances of the word 'green' are different within the utterances of 'it is green'. But then the contents of the utterances of 'green' within the utterances of 'the leaf is green' must be different as well. *Q.E.D.*

Let me introduce some abbreviations and restate the argument to make its structure clearer:

- U_A: A's utterance of 'the leaf is green'
- u_A: A's utterance of 'green' within U_A
- U_B: B's utterance of 'the leaf is green'
- u_B: B's utterance of 'green' within U_B

- U'_A: A's utterance of 'it is green' following U_A
- u'_A: A's utterance of 'green' within U'_A
- U'_B: B's utterance of 'it is green' following U_B
- u'_B: B's utterance of 'green' within U'_B

I will use 'C(v)' to refer to the content of v. Travis's claim is that C(U₁) ≠ C(U₂) even though C(u₁) = C(u₂), which would be a counterexample to the principle of

³¹ This is not to say that he said the same thing on both occasions. Arguably, in uttering 'the leaf is green' A said of the painted leaf that it is a leaf, but in uttering 'this is green' he said no such thing.

compositionality of speech-act content. I granted the first claim and argued against the second. My argument went as follows:

$C(u_1) = C(u'_1)$

$C(u_2) = C(u'_2)$

$C(u'_1) \neq C(u'_2)$

Therefore, $C(u_1) \neq C(u_2)$

The first two premises were based on the intuition that in restating earlier statements the contents of restated words remain the same. The third premise is supported by the following sub-argument:

A and B said different things

A and B were talking about the same leaf

Therefore, A and B said different things about the same leaf

Whatever A and B said about the leaf was said in calling it 'green'

Therefore, A and B said different things about the same leaf in calling it 'green'

In order to facilitate the intuitive assessment of the premises I have kept the argument as close the vernacular as possible. But I think we do have established technical terms for the speech-acts involved here. To say that A and B said different things is to say that they *asserted* different things, to say that they were talking about the same leaf is to say that they *referred* to the same leaf, and to say that they said different things about the leaf is to say that they *predicated* different things of it. The conclusion of the sub-argument is that A and B predicated different things of the leaf in calling it as 'green'.

It is clear that the same style of argument applies in a number of similar cases – utterances of 'this flea is small', 'the table is flat', or 'the children are ready'. Take the last example. One says one thing in uttering 'the children are ready' when they are standing next to the pool preparing to jump, and another thing when they are sitting down at their desks to take their math exams. In response to the question 'what did you say the children are?' one would be prepared in both cases to respond 'they are ready' and in doing so one would use the word 'ready' with the same content as before. In making the

two utterances of ‘they are ready’ one would be referring to the same children, predicating different things of them. But whatever is said about those children in making these utterances is said in calling them ‘ready’. So the contents of the two utterances of ‘ready’ within the utterances of ‘they are ready’ are different, and hence, the contents of the two utterances of ‘ready’ within the utterances of ‘the children are ready’ are different too.

There are other cases that require some extra maneuvering. Take the case of the possessive construction. Different things can be said in uttering ‘Jim’s coat is black’ – one might say, for example, that the coat Jim owns is black, or that the coat he is sitting on is black, or that the coat he stole is black, etc. What we want to show is that in these cases the utterances of the possessive morpheme ‘-s’ also have different contents. The problem is that there are no subject-predicate sentences where the possessive morpheme is predicated of something, and so we do not have a sentence that could play the role of ‘it is green’ or ‘they are ready’ in the previous cases. But the difficulty is not insurmountable.

The key is to replace the original example with ‘this is Jim’s coat’ uttered while pointing at a particular coat. Utterances of this sentence display the same variability in content as utterances of ‘Jim’s coat is black’, so it would be quite implausible to think that the content of utterances of ‘-s’ varies within utterances of one but not within utterances of the other. So consider A’s utterance of ‘this is Jim’s coat’ whereby A says that the object he is pointing at is one Jim owns and B’s utterance of the same sentence whereby B says that the object she is pointing at (let’s suppose this is the same object) is one Jim was previously sitting on. The next step is to consider the question ‘whose coat did you say this is?’ to which the speakers respond by uttering ‘it is his’ pointing at Jim. I take it that the speakers’ utterances of ‘his’ count as restatements of their earlier utterances of ‘-s’, and hence, that (setting aside the indication of gender) these utterances have the same content. Now we can continue the argument as before reaching the conclusion that the two utterances of ‘his’ within ‘it is his’ have different contents, and hence, that the two utterances of ‘-s’ within ‘this is Jim’s coat’ have different contents too.

Another example where extra maneuvering is required is that of quantificational domains. If A and B coming out of their respective classrooms utter ‘most students passed’, they say different things – A that every student in A’s class passed, B that every student in B’s did. Here we need to replace the original example with a partitive construction – ‘most of the student passed’ – and consider the question ‘who are these people of whom you said most passed?’ In answering the question by uttering ‘they are the students’ A and B said different things (each referring to the students in their respective classes), and uttered ‘the students’ with the same content as before. The argument goes as before and leads to the conclusion that the contents of utterances of ‘the students’ within ‘most of the students’ differ. This, I think makes it overwhelmingly plausible that the contents of utterances of ‘students’ within ‘most students’ differ too.³²

I gave a general template for showing that whenever a sentence s is used to say different things it has a constituent expression e that is also used to say different things. The key is to find an appropriate sentence s' (i) whose subject is a pronoun whose reference is fixed across the different utterances and whose predicate is e (or a closely related expression – ‘his’ instead of ‘-s’, ‘the students’ instead of ‘students’, etc.), (ii) whose utterances vary in content in a way that is analogous to the variation in content in utterances of s . There is no *a priori* guarantee that such an s' can always be found. Still, the examples above strongly suggest that the search will not be in vain. The fact that in these diverse cases of differences in the speech-act content of sentential utterances we could find differences in the speech-act content of constituent utterances makes a strong empirical case that (C_{sc}) holds.^{33,34}

³² It might seem that the reasoning remains inconclusive due to an unnecessary detour. But this is not true. Suppose the question ‘who are these people of whom you said most passed?’ is raised after utterances of ‘most students passed’ and suppose A and B answer by uttering ‘they are students’. It is true that A and B said different things and also that the contents of their utterances of ‘students’ were the same on the two occasions. But they were talking about different students, so the rest of the argument is blocked.

³³ As it stands, the template can only discard alleged counterexamples to (C_{sc}) involving *sentences* that are used to assert different things. All the alleged counterexamples in the literature are of this sort. I don’t think this is an accident – it is hard to argue that a complex expression can be used to make speech-acts with different contents unless those speech-acts are illocutionary acts, and if they are one might as well use sentences to express the relevant content.

5. Conclusion

Is linguistic content compositionally determined? Traditional arguments – productivity and systematicity – do not speak to the matter. But there is a straightforward argument for the compositionality of expression content from the compositionality of speech-act content, and there is a template to respond to a wide array of challenges to the compositionality of speech-act content. This warrants the cautious belief that our semantic theories are not as misguided as Fodor thinks – linguistic content, whether it be the content of an expressions within the context of its utterance or the content of a speech-act performed by uttering an expression, is determined bottom-up.³⁵

³⁴ There is an important putative counterexample to the principle of compositionality of speech-act content I did not discuss. According to Fine (2007), ‘Cicero is Cicero’ and ‘Cicero is Tully’ differ in content, even though proper names are directly referential. If so, (C_{cc}) fails for English and – assuming (L₁) – literal uses of these sentences yield a counterexample to (C_{sc}) too. However, I think an application of the general argumentative strategy pursued above suggests strongly that *if* we accept that speakers who utter these sentences assert different things *then* we should also accept that in uttering ‘is Cicero’ and ‘is Tully’ they predicate different things of the famous Roman orator. I intend to discuss Fine’s semantic relationism further elsewhere.

³⁵ I thank audiences at Columbia, Rutgers, Santa Cruz, the 2007 Pacific Division Meeting of the APA in San Francisco, and the 15th Annual Meeting of the European Society for Philosophy and Psychology in Geneva, where I had the opportunity to present earlier versions of this paper. Special thanks to Sandra Chung, Chris Gauker, Tamar Szabó Gendler, Martin Jonsson, Peter Lasersohn, Ernie Lepore, John MacFarlane, Josep Macia, Christopher Peacocke, and Daniel Rothschild for discussion.

References

- Bach, K. (1994) 'Semantic slack: what is said and more.' In S. Tsohatzidis ed., *Foundations of Speech-Act Theory*, London: Routledge, pp. 267 – 91.
- Bach, K. (2001) 'Speaking loosely: sentence nonliterality.' In Peter French and Howard Wettstein eds., *Midwest Studies in Philosophy* **25: Figurative Language. Oxford: Blackwell, pp. 249-263**
- Cappelen, H. and E. Lepore (2005) *Insensitive Semantics*. Oxford: Blackwell.
- Carnap, R. (1947) *Meaning and Necessity*. Chicago: University of Chicago Press.
- Carston, Robyn (2002) *Thoughts and Utterances*. Oxford: Blackwell.
- Fine, K. (2007) *Semantic Realism*. Oxford: Blackwell.
- Fodor, J. (2001) 'Language, thought and compositionality.' *Mind and Language* **16**: 1–15.
- Fodor, J. and E. Lepore (1991) 'Why meaning probably isn't conceptual role.' Reprinted in *The Compositionality Papers*. Oxford: Blackwell, 2002, pp. 9 – 26.
- Higginbotham, J. (1988) 'Contexts, models and meanings: a note on the data of semantics.' In R. Kempson ed., *Mental Representations: The Interface between Language and Reality*. Cambridge: Cambridge University Press, pp. 29 – 48.
- Horwich, P. (1998) *Meaning*. Oxford: Clarendon Press.
- Janssen, T. M. V. (1983) *Foundations and Applications of Montague Grammar*. Amsterdam: Mathematisch Centrum.
- Janssen, T. M. V. (2001) 'Frege, contextuality and compositionality.' *Journal of Logic, Language, and Information* **10**: 115 – 136.
- Kaplan, D. (1977) 'Demonstratives. An Essay on the Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals.' Reprinted in J. Perry and H. K. Wettstein eds., *Themes from Kaplan*. Oxford: Oxford University Press, 1989, pp. 481 – 565.
- King, J. and J. Stanley (2005) 'Semantics, pragmatics, and the role of semantic content.' In Z. G. Szabó ed., *Semantics vs. Pragmatics*. Oxford: Oxford University Press, pp. 111 – 64.
- Lahav, R. (1989) 'Against compositionality: the case of adjectives.' *Philosophical Studies* **57**: 261–279.
- Levinson, S. (2000) *Presumptive Meanings*. Cambridge, MA: MIT Press.
- Lewis, D. (1970) 'General semantics.' Reprinted in *Philosophical Papers. Vol. 1*. Oxford: Oxford University Press, pp. 189 – 229.
- MacFarlane, J. (forthcoming) 'Nonindexical contextualism.' *Synthese*.
- Pelletier, F. J. (2001) 'Did Frege believe in Frege's Principle?' *Journal of Logic, Language, and Information* **10**: 87–114.
- Pelletier, F. J. (2003) 'Context-dependence and compositionality.' *Mind and Language* **18**: 148–161.
- Perry, J. (1986) 'Thought without representation.' In *Proceedings of the Aristotelian Society. Suppl. Vol. 60*: 263 - 83. London: Aristotelian Society Publications.
- Recanati, F. (1989) 'The pragmatics of what is said.' *Mind and Language* **4**: 295 – 329.
- Searle, J. R. (1969) *Speech Acts*. Cambridge: Cambridge University Press.

- Searle, J. R. (1978) 'Literal meaning.' *Erkenntnis* **13**: 207 – 224.
- Soames, S. (2005) 'Naming and asserting.' In Z. G. Szabó ed., *Semantics vs. Pragmatics*. Oxford: Oxford University Press, pp. 356 – 82.
- Sperber, D. and D. Wilson (1986) *Relevance: Communication and Cognition*. Cambridge, MA: Harvard University Press.
- Stainton, R. (2006) *Words and Thoughts*. Oxford: Oxford University Press.
- Stalnaker, Robert (2002) 'Common ground.' *Linguistics and Philosophy* **25**: 701 – 721.
- Stanley, J. (2000) 'Context and logical form.' *Linguistics and Philosophy* **23**: 391 – 434.
- Stanley, J. and Z. G. Szabó (2000) 'On Quantifier Domain Restriction' *Mind and Language* **15**: 219 - 61.
- Szabó, Z. G. (2000a) 'Compositionality as supervenience.' *Linguistics and Philosophy* **23**: 475–505.
- Szabó, Z. G. (2000b) 'Adjectives in context.' In I. Kenesei and R., Harnish eds., *Perspectives on Semantics, Pragmatics, and Discourse*. Amsterdam: John Benjamins, pp. 119–146.
- Szabó, Z. G. (2000c) *Problems of Compositionality*. New York: Garland.
- Travis, C. (1985) 'On what is strictly speaking true.' *Canadian Journal of Philosophy* **15**: 187 – 229.
- Travis, C. (1994) 'On Constraints of Generality.' In *Proceedings of the Aristotelian Society*. New Series **44**: 165 - 88. London: Aristotelian Society Publications.
- Walton, K. (1990) *Mimesis as Make-Belief*. Cambridge, MA: Harvard University Press.
- Zadrozny, W. (1994) 'From Compositional to Systematic Semantics.' *Linguistics and Philosophy*, **17**: 329 – 342.