10

Context, Compositionality, and Brevity*

FRANCIS JEFFRY PELLETIER

10.1. Brevity—why bother?

Intuitively, the title of this section offers two rather different questions: one is to query why we should be brief in our speech and the other is to question whether it is worthwhile (or even possible) to try to give an *account* of the fact that people exhibit brevity in speech. The title, in fact, *exhibits* a kind of brevity in posing these two different questions in just a few words. To answer the first of the two questions about *this* example of brevity, we might say such things as: we wanted to highlight the two aspects but they wouldn't both fit as a section title, or that we wished to engage those who are only interested in one aspect into a discussion of both questions, and so on. In fact, perhaps we didn't ourselves even notice, as we were composing the title, that it had these two aspects. And so on.

   The second question is the focus of the present book as a whole: it is *presumed* that there are a number of different ways that brevity occurs in speech—for example, by ellipsis, by appeal to common ground or social practice, by recognition of others' communicative intentions, by forming joint communicative intentions, by accommodation, by uttering sub-sentential phrases, and by various other methods—and the goal is both to characterize these methods clearly and to give a theoretical account of how, despite the appearance that important information has been 'left out', communication nonetheless can succeed.

   Often when theorists are answering the second question, they implicitly (or sometimes explicitly) also give an answer to the first question. For example, it might be said that fully to articulate all (or even some) of the common ground

would be immensely inefficient—and indeed, might convey further information which was not at all desired to be communicated by the very act of trying to explicitly exhibit it. So, never to engage in *any* form of ellipsis or brevity would make speech intolerably tedious. Never to engage in *any* form of 'mind reading' would be to place oneself outside the social society of one's interlocutors.

These sorts of reasons were parodied by Jonathan Swift:[1]

We next went to the School of Languages, where three Professors sat in Consultation upon improving that of their own country.

The first Project was to shorten Discourse by cutting Polysyllables into one, and leaving out Verbs and Participles, because in reality all things imaginable are but Nouns.

The other, was a Scheme for entirely abolishing all Words whatsoever; and this was urged as a great Advantage in Point of Health as well as Brevity. For it is plain, that every Word we speak is in some Degree a Diminution of our Lungs by Corrosion, and consequently contributes to the shortening of our Lives. An Expedient was therefore offered, that since Words are only Names for Things, it would be more convenient for all Men to carry about them, such Things as were necessary to express the particular Business they are to discourse on . . .

Another great Advantage proposed by this Invention, was that it would serve as a Universal Language to be understood in all civilized Nations, whose Goods and Utensils are generally of the same kind, or nearly resembling, so that their Uses might easily be comprehended. And thus Ambassadors would be qualified to treat with foreign Princes or Ministers of State to whose Tongues they were utter Strangers.

On the opposite hand, considerations of compositionality have led some to find that brevity has its own shortcomings. If one intends to convey a message—say, the message JIMMY HAS JUST FALLEN INTO THE WELL—one needs to produce enough language so that the hearer understands. One cannot reasonably omit that part of the language that corresponds with INTO THE WELL, nor that with FALLEN, nor that with JIMMY. One cannot even omit something corresponding to INTO, without the danger that the hearer simply will reject the utterance as ungrammatical. (Of course, none of this is to say that the linguistic representation of that message must contain a lexical item for each part of that message: rather, it is that the message as a whole must be conveyed. Some linguistic methods of expression—English, for example—may do it with a word that corresponds to each part of the message.[2] But other languages might do it with a case marking, or with a position in a sentence, or with prosody, clitics, particles, or with a variety of other linguistic methods.)

---

[1] *Gulliver's Travels*, Book III, Chapter V 'The Author Permitted to See the Grand Academy of Lagado', midway through the chapter.
[2] Not surprising, since this article is written in English, and so I have expressed the message in (small-caps) English. But even if messages were written in neural-speak, then there might still be some language that maps that onto a word-for-word transcription.

So, it is obvious that the medium in which the desired message is going to be encoded must contain sufficient structure and content that the hearer can determine the message. And the most straightforward way to do that is to encode the message in a compositional manner. This claim—that the most straightforward way to linguistically encode a message is to employ a compositional semantics—is an instance of the so-called 'Argument from Productivity', which is one of the three considerations that are used to support semantic compositionality.

The Argument from Productivity/Creativity: People can produce a sentence that they have never heard, and know that it is not only a part of their language but also what it means, and that it expresses what they are trying to communicate. The only explanation that has ever been offered for this is that language is semantically compositional.

The Argument from Understandability: People can hear a sentence that is entirely new to them, and as long as it contains words and syntactic structures that they are familiar with, they will understand the message that the sentence conveys. The only explanation that has been offered for this is that language is semantically compositional.

The Argument from Learnability: Children learn the meanings of an unbounded number of sentences in a finite time. The only explanation that has been offered for this is that language is semantically compositional.

Together these arguments might be called Arguments from the Infinity of Language (and the finiteness of humans). They pose a challenge to show how an 'unbounded' ability can arise from 'finite' creatures. All make the claim that semantic compositionality provides a way; and that there is no other way known. Of course, this is an argument to the best explanation, not a deductive argument; and so the conclusion may be false. But the Arguments from the Infinity of Language have been thought to be quite persuasive, especially in the absence of any credible alternative mechanism.

10.2. Context and compositionality[3]

Context can, however, ease the restriction that the message must contain 'full information'. If, for example, participants in a conversation have become aware, for whatever reason and in whatever way, that the topic of the conversation has been the old well in the south 40, then saying 'Jimmy has just fallen into it' would do the trick. Similarly— *mutatis mutandis*—were the topic of the context to be Jimmy, then saying, 'Just fell into

---

[3] I take for granted here—what is a dubious assumption—that we are all on the same page when it comes to the notions of 'context' and 'compositionality'. I will be a bit more definite later about the notions, but one should consult some standard works to see the range of interpretations these notions have given rise to. One can consult many of the articles in Werning et al. (2012), such as Recanati (2012); Pelletier (2012); Westerståhl (2012) for discussions of context and compositionally. On the topic of context in particular, see Recanati (2004). One can also find an annotated bibliography of 150 works on compositionality in Pelletier (2011).

the well' would carry the desired information. And a variety of other contextual features could allow the message to be made briefer in a number of other ways. I intend to discuss a certain class of these contextual features to see whether, and if so, how, they can be compatible with compositionality. For after all, communication takes place in a context, and we have already indicated that there is no known mechanism other than compositionality that allows for successful communication that obeys the dictates of the Arguments from the Infinity of Language. So it seems that there *must* be a way to make contextual features compatible with compositionality.

10.2.1. *There are contexts, and then there are contexts*

The literature makes a major divide between two different types of context to be considered in evaluating some piece of communication: the 'linguistic context' and the 'situational context'. As the names suggest, the difference depends on whether the context is given by means of words/sentences (in the preceding or nearby vicinity) or if it is the non-linguistic setting of the communication. The difference is clear in paradigmatic cases, but cases can blend into one another.

   While admitting that the distinction between the two types of context can be blurry, we might still focus on clear examples of linguistic vs. situational cases. On the linguistic side, we see cases where certain types of 'expectations' are developed by (for example) explicit assertions to the effect that utterances in the immediate future of the conversation should be interpreted in a certain way. For example,

(10.1)   Ok, we're done talking about non-poisonous snakes; we now move on to the poisonous snakes. The Australian ones are particularly notable.

In this discourse, the interpretation of *ones* is determined by the preceding sentence, and must mean poisonous snakes.[4] Generally speaking, this way of looking at context as being set linguistically suggests that a discourse starts out with some presumed-by-the-conversational-participants background 'model' and then updates the model with new information in a sentence-by-sentence manner. Each new sentence is evaluated for truth in terms of the model as thus far developed, and in turn, the sentence also updates the model further. Even though the details of how this evaluation-and-update model works are rather mysterious, I shall employ it in an impressionistic way, so as to show how I see the way compositionality and context can interact. The details are important, but I also believe their final analysis will not change the general thrust of what I say. (The interested reader might consult Eijck and Visser (2010), as well as the works inspired by Robert Stalnaker's publications, as

---

   [4] And before one starts thinking, 'No, no: we don't need any fancy interpretation-in-a-model, all we need is to expand the notion of anaphora so that it can cross sentence boundaries', note that the same principle holds were the final sentence 'The Australians are particularly notable'. In such a continuation we still need to interpret 'Australians' as Australian snakes.

reprinted in Stalnaker (1999) and also Isard (1974), Lewis (1979), Stucky (1989), and many of the contributors to Akman et al. (2001)).[5]

Yet a different aspect of the effect of context on interpretation occurs when the *non*-linguistic context fixes or alters the meaning of terms and larger pieces of discourse. Of course, the presence of salient entities (and perhaps their gender) in a physical situation where a sentence with a pronoun is being uttered can (help to) fix the referent of the pronoun. Similar considerations work also for ('incomplete') definite descriptions, such as 'The man', whose identity in a context may be determined by there being only one salient adult male.

Slightly further away from this sort of assignment-by-salience are such very ordinary cases as (from Rumelhart 1979: 78):

(10.2)   The policeman stopped the car.

which, depending on a variety of features of a context, might mean that the policeman ordered the driver to stop the car and succeeded, or might mean that the policeman *was* the driver and he pushed on the brake pedal, or might mean that the policeman set up a physical roadblock that forced the car to crash, or might mean that the policeman was in fact Superman and stopped the car by lifting it off the ground, etc. According to some, these are in fact different 'occasion' meanings of (10.2).

Yet further away are examples of the sort favoured by Charles Travis (1996; 1997), where the occupations and interests of the speakers play an important role in determining the truth-value of sentences such as:

(10.3)   These leaves are green.

which, when said of painted leaves might be thought true by painters or crime-scene investigators, but false by botanists, for example. In general, it is thought, the physical aspects of a context-of-utterance, the identities/occupations/etc of the conversational participants, and numerous other factors, can and do play roles in determining the meanings of words/phrases/sentences that are uttered in those contexts.

10.2.2. *Speaker meaning and cccasion meaning*

'Speaker meaning' centrally includes the idea that the speaker intends that the hearer will come to have a belief (or other response), and will do so on the basis of the hearer's recognition that the speaker wanted him to respond in that way. Put this

---

[5]   One special case of linguistic context occurs when some part of the sentence currently being evaluated by means of some already-established model itself updates the model in such a way as to affect the further evaluation of the rest of the sentence. While many such cases can be naturally handled by the general method suggested above, there are a number of recalcitrant instances that have been seen as challenges to the compositionality methodology. (Some of these are outlined in Pelletier (1994); a particularly nice example derives from Higginbotham (1986)). As interesting as these are, we will not pause over them in this chapter.

way, a speaker can employ aspects of the situation in which he is making his utterance in order to get the hearer to have the right response (for the right reasons).

In the current discussion, this means that there are two rather different conceptions of 'meaning in context', both of which could be (and have been) called 'occasion meaning'. One conception—the one that I have implicitly been suggesting—is where the situational features themselves induce an alteration from the literal meaning. But the other conception is equally plausible: that the speaker is employing these situational features in constructing his speaker meaning. (See Bach 2012, 2013, for a defence of this as the 'real' notion of occasion meaning). We might call these two senses the 'objective' sense of occasion meaning versus the 'intentional sense' of occasion meaning, respectively.

I am going to indifferently call both of these notions 'occasion meaning', even though in both the metaphysics of the two theories and in their views about where to look for meanings, the two are really quite different. I justify my indifference between the two conceptions by thinking of occasion meaning in the objective sense as what is common to all, or most, or the typical, or the ordinary occasions in which a speaker is employing features of the context in the intentional sense of occasion meaning. That is, the objective sense is what a reasonable speaker should expect a reasonable hearer to be able to figure out was the speaker's meaning in that situation.

10.2.3. *Semantic compositionality*

Semantic compositionality is the principle that the meaning of a whole depends on the meanings of its parts and the way the parts are combined. This is a pretty vapid statement unless one sharpens various concepts used in this statement of the principle. By 'whole' is normally meant some syntactically complex item of language; by 'its parts' is meant the syntactic parts of this whole; and by 'way combined' is meant the syntactic manners of combination.[6] A crucial notion in the principle is 'depends on'. This has been given many interpretations but we will understand it as implying 'is a function of' (in the mathematical sense). Thus the principle becomes:

The meaning of a syntactically complex item is a function (in the mathematical sense) of the meanings of its syntactic parts and the semantic effect of the syntactic manner in which these parts are combined.

It should be noted that employing the concept of a function here entails that the meanings of complexes may contain 'items' and 'information' that are not found in their syntactic parts; for a function can introduce 'new material' into the meaning of

---

[6] The 'and the way they are combined' condition is what allows us to distinguish the meanings of *Democrats distrust Republicans* and *Republicans distrust Democrats*, for example, even though they are both 'made up of the same parts'.

the complex, so long as it does this uniformly whenever the sub-expressions are of the same syntactic types.

Of course, this has not explained what is intended by 'meaning' in the statement of the principle, and because of the goals to be described below, I will not further define the notion at this time. Eventually, I want to include some of the more common notions of meaning in the literature, and show how they can be given a compositional treatment. In doing this we will describe two more-or-less standard areas where 'meaning' is used, showing their relevant differences that transcend whatever notion of meaning is being employed.

These two more-or-less standard areas are generated by a distinction between what are called 'literal meaning' and 'occasion meaning'.[7] The intuitive idea is that literal meaning encodes what words and syntactic constructions contribute merely as their role in the language, whereas occasion meaning reflects what the words and constructions have evolved to contribute in the course of a conversation that takes place in a certain situation. In the framework cited earlier, where a discourse is envisaged as a series of updates to a starting context, the literal/occasion distinction can be seen as the difference between a starting context and any of the further contexts. Of course, in the usual case the starting context *does* contain special information: the participants can detect their physical surroundings, they know at least visual information about their interlocutors, and so on. Thus there is no truly 'neutral context'. But as an idealization, if the initial context contains no special information we might say that the language used employs the literal meaning. Or, one may try to define the 'neutral context' as what is common to the initial contexts of all conversational situations. We will not pause over the exact account of such a notion, since the distinction as we will employ it will always be clear.

The general question under consideration is whether compositionality can be maintained even in the face of occasion meaning. But there are a number of different types of objections to this claim, and it is worthwhile to first discuss some of the more general sorts before moving on to what I think has often been ignored about literal and occasion meanings.

10.2.4. *Some anti-compositionality-of-context claims (and what's wrong with them)*

10.2.4.1. *Generality*   One very general objection to semantic compositionality occurs when a context can 'change the meaning' of a term. For example, 'red' is said to change its meaning depending on whether the context is about hair, wine, fruit (and even species of fruit), automobiles, blood, etc. Thus the colour being described in 'red apple', 'red wine', 'red hair', 'red watermelon', 'red skin', 'red blood', etc, differs. And if one thinks of the meaning of 'red' in terms of its 'recognition features' then the

---

[7] Recall that I use this term in my 'indifferent' way.

meaning of 'red' has changed between different contexts and is thus a violation of compositionality.[8]

John Searle (1980: 223) says much the same about 'cut' (*cut the lawn* versus *cut the cake*), saying 'that word ['cut'] is interpreted differently in each sentence.' Recanati (2012: 177) remarks that the same is true of 'likes':

He *likes* my sister. He *likes* roasted pork.... The first sentence talks about 'affective' liking and the second about 'culinary' liking. There is as much difference between the two kinds of state as there is between the processes of cutting involving grass and cakes.

And L. Jonathan Cohen (1986: 228) says:

Most students here *drop* geography in their final year. Most students here *drop* geography lectures in their final year. Most students here *drop* geography lectures reading assignments in their final year. Most students here *drop* geography lectures reading assignments library-fees in their final year ... [T]here seems no predictable end to the variety of expressions that we can put meaningfully after 'drop', so as to impose a series of different meanings on the latter word.

Taking this line of thought seriously, Renate Bartsch (1994) denies that there is any literal meaning at all, matched by similar argumentation and conclusions in Charles Travis (1996, 1997).

It often seems to me that such critics of literal meaning are off the mark. For one thing, it has yet to be shown that a meaning of a word or phrase or sentence (whether literal or occasion meaning) is not best represented as a *set* of more specific or particular meanings. If anti-compositionality-of-context is just the claim that different contexts yield different meanings to the same syntactic item, then why not just make the literal meaning be the set of all these specific-to-the-context meanings? Furthermore, if 'cut' really changed its meaning between a context where we were discussing grass mowing and one where we were discussing dessert serving etiquette, then it ought to be impossible to order someone, in the former context, to cut the lawn with a knife. But this is quite obviously possible; and according to classic movies concerning US Marines boot camps, one can clean bathroom floors with a toothbrush. The fact that a method is not standard in a given context does in no way curtail its being permissible in that context; and as a corollary, does in no way prevent it from being an aspect of the literal meaning of the word/phrase/sentence.

One is in fact quite tempted to claim that these critics simply don't understand what generality is. If I say, 'John is meeting somebody on Wednesday', I have *not* meant *anything* about the gender of the referent of 'somebody'. And that remains true even if I *knew* that it was a woman he is meeting; for, information about this is

---

[8] One can also make this be a case of an intra-sentential context shift, since the truth-conditions of the colour attribution 'obviously' depend on the noun being modified. And since the noun is not a syntactic subpart of the adjective, this becomes a violation of compositionality even *within* one sentence. But I said earlier that I would not consider these sorts of intra-sentential contexts.

not at all anything I meant by uttering the sentence. The word 'somebody' is simply *general* about (or, if you prefer, *unspecified* for) gender. It would be totally *silly and absurd* to insist that the context must somehow generate the meaning WOMAN out of the utterance of 'somebody', regardless of whether or not I have that information myself.

Of course, sometimes the speaker *does* want the hearer to infer that the 'someone' is a woman. And there can sometimes be situational features that make it possible for any ordinary speaker to believe that an ordinary hearer will be able to infer that he (the speaker) is trying to convey this information. But it is important to note that the mere desire on the part of the speaker that this should happen is not sufficient for it to be part of the speaker meaning (and hence not a part of the occasion meaning in either the objective or the intentional sense). Not all desires that a hearer should come to believe something are *intentions* of the relevant type. For, there just are certain intentions that a speaker just can't have: it is not possible to have an intention that you *know* cannot be fulfilled. Nor even one for which you firmly believe that there is no reason whatsoever to think that it will be fulfilled. Nor is it coherent for you to have an intention that someone else should do something that you know to be impossible for them to do.

If these are not possible intentions for you to have, then you can *not* have the intention that your hearer should believe that the person being met was female by virtue of the listener recognizing your intention, if all the evidence your hearer has is that you say, 'John is meeting *somebody* . . .'. And from hearing that alone, it is not possible for your hearer to recognize your so-called intention that they come to believe that the person being met is female (just on the basis of hearing those words). But surely you know this. And therefore your alleged 'intention' in fact was not really an intention at all. And hence is not your speaker meaning either. So it can't be the occasion meaning—unless the situation already had other features, such as that perhaps the conversation was already about the various women in John's life. Perhaps then it is possible that you could intend the hearer to infer that the 'somebody' was female. But then, that doesn't challenge compositionality. So none of that is germane to the meaning, on that occasion, by that utterance. Nor of course is it relevant to the *literal* meaning of the sentence and the word 'somebody'.

As I said above, it seems that some people just don't know what *generality* is!

And now: is there *any reason whatsoever* to think that the case with 'cut' or 'like' or 'red' is different in any way at all from the case of gender? I think not. None of these is a challenge for compositionality. They only gain any traction by insisting that meaning is totally specific and never general.[9]

---

[9] The consideration of generality discussed in this section does not, of course, touch on the Travis examples in which it is the occupation of the conversationalists that affects the truth value of the utterance. Under the heading of 'pragmatic groundedness' I later intend to deal with such cases.

10.2.4.2. *'Good', 'large', 'small', and the like*   Another set of anti-compositionality considerations comes from the existence of 'relative adjectives' whose scale apparently depends on the noun being modified. The classic examples involve 'large' and 'small', although as long ago as Aristotle philosophers discussed 'good' and 'healthy' in much the same way.

(10.4)   Mickey is a big mouse; All mice are animals; but Mickey is not a big animal.

(10.5)   Sarah is a bad violinist; All violinists are people; but Sarah is not a bad person.

A tradition in formal semantics, probably starting with Terence Parsons (1968, 1970) and Romane Clark (1970), and gaining wide acceptance by way of Kamp (1975) together with some of the suggestions in Montague (1970), is to treat these kinds of adjectives as somehow 'dependent upon the noun being modified'. For example, the adjective 'big' might have a variable in its meaning that takes the noun being modified and produces a complex meaning:

(10.6)   big $\Rightarrow \lambda X \lambda x[X(x) \wedge (\text{BIG FOR AN } X)(x)]$[10]

One could apply the suggestion to other terms, for instance to verbs like 'cut' that depend for a meaning given by the direct object:

(10.7)   cut $\Rightarrow \lambda X \lambda x \lambda y[X(y) \wedge (\text{CUT IN THE MANNER OF } X)(x, y)]$

In these cases the Xs get filled in by the (meanings of the) relevant nouns (e.g., 'mouse', 'elephant', for 'big'; and 'grass', 'cake' for 'cut'). Once one sees how to do this in the simple cases, it becomes tempting to do it for *every* adjective—saying in effect that no adjective corresponds to an independent set of objects to which it 'applies in the abstract', but rather it is *always* dependent for a noun to pick out some set of things and then for the adjective to narrow down that set in the appropriate way. You first pick out the set of mice, then you find the big ones of them; you first pick out the set of animals, then you find the big ones of them. You first pick out the set of violinists, then you find the bad ones of them. Even with colour terms: You first pick out the set of apples, then you find the red ones of them; you first pick out the set of (human) hairs, then you find the red ones of them.

In the case of transitive verbs, you first pick out the object-kind under consideration: perhaps Grass. Then you look at the activities of cutting it. Or, you pick out the object-kind, Cake, and then you look at the activity of cutting it. Just as with the adjectives, where there is no presumption that the large mice will be cases of large animals, with verbs there is no presumption that the cuttings-of-grass will be cases of cuttings-of-cake. (And even if all grass is foliage, it doesn't follow that all

---

[10] Put this way for illustration only. A term like 'big' cannot be *defined* by 'big for an X', on pain of circularity. Rather, the interpretation is to pick out a subset of the Xs that are independently given as the relevant ones. Thanks to Kent Bach for insisting on this point.

cuttings of grass are cuttings of foliage.) Just as there is nothing that corresponds to the adjective 'large' without the intervention of some noun being modified,[11] so too there is nothing that corresponds to the verb 'cut' without the intervention of some object.

Once one sees the meaning of adjectives and verbs in this manner, the anti-compositional force of these types of examples loses its strength. The 'generality' considerations of the preceding subsection and the 'complex modification' considerations of the present discussion are each aimed at the same targets, and it might seem that we need not have both sets of considerations in order to defend compositionality in the face of these examples. My own view is that the generality considerations work best in defending against some of the examples (eg, the Searle 'cut'-style cases) whereas the modification considerations work best in defending against the adjectival examples like 'red'. But my goal here is only to show that the case against compositionality that is brought up by such examples is not so clear-cut as some scholars have thought. We will consider the role of context more in the next few sections.

I should also remark here that, as noted already by Recanati (2012), the idea of allowing the context to develop the appropriate occasion meaning for the relevant nouns first, and then to apply the *big*, *good*, *short*, etc, adjectives to the occasion meaning of the right noun, neatly avoids an objection raised by Irene Heim and Angelika Kratzer (1998) and also an earlier one of Searle's (1980).

Imagine we had first introduced a scenario populated with an army of monsters like King Kong. We might then have said something like 'Jumbo doesn't have a chance; he's only a small elephant', and this could have been true even if Jumbo were as large or even larger than most other elephants (Heim and Kratzer 1998: 71).

It is easy to imagine circumstances in which 'cut' in 'cut the grass' would have the same interpretation it has in 'cut the cake', even though none of the semantic contents of the words has changed. Suppose you and I run a sod farm where we sell strips of grass turf to people who want a lawn in a hurry.... Suppose I say to you, 'Cut half an acre of grass for this customer'; I might mean not that you should *mow* it, but that you should slice it into strips as you could cut a cake or a loaf of bread (Searle 1980: 224–5).

The idea here is that a small elephant is not necessarily an elephant that is small for an elephant. In this example, it is supposed to be small-for-a-monster, as given by the context. And 'cut the grass' does not mean cut in the manner of grass-cutting but rather some other manner given by the context.

---

[11] Well, maybe the union of all of the *large Xs*, for every noun *X*? But then things would fall into *both* the extension of 'large' and of 'small'. Probably also into 'not-large', under plausible assumptions concerning the relation of 'small' and 'not-large'.

10.3. Pro-compositionality-and-context suggestions

Some contextualist objections to compositionality seem to operate by deliberately confusing the standard and occasion meanings. Consider, for example, the claim that the meaning of:

(10.8)   That new prospect is very big

can't be compositionally derived from that of its parts because we first need to know what 'category' the new prospect is in. Is the prospect a basketball draftee? Or may be a business opportunity? Or some other thing?...Without that information, we cannot compute the meaning of (10.8).

And it is not just the individual words which get their meaning assigned by context, but even the semantic effect of certain syntactic constructions gets assigned by context. Consider, for example, non-restrictive relative clauses:

(10.9)   Sandy, whose paper was rejected by *Linguistics and Philosophy*, began to rant.

In some contexts—perhaps the most normal ones—the semantic effect of this non-restrictive relative clause is BECAUSE. But other times—again, depending on context—it can have the semantic effect of DESPITE or AND or ALTHOUGH, among yet further things.

Note, however, that these various meanings—the BASKETBALL PLAYER or vs. BUSINESS DEAL; the BECAUSE or vs. DESPITE, etc—are all involved with describing the *occasion* meanings of the sentence. But: if you consider the *occasion* meaning of the entire sentence, then compositionality merely says that it can be described by the *occasion* meanings of its parts and the *occasion* semantic effect of the ways those parts were syntactically combined...and these meanings arguably *do* have the relevant information about what 'category' the new prospect is in. On the other hand, if you consider the *literal* meaning of the entire sentence, then compositionality claims merely that *it* can be described by the *literal* meanings of its parts and the *literal* semantic effect of the way those parts were syntactically combined—but naturally, you won't know what 'category' the new prospect is in, since that's not a part of the literal meaning, but only of the occasion meanings. The only difficulty for compositionality comes when one tries to generate occasion meanings of wholes out of the literal meanings of parts. (Or vice versa, but nobody tries to do that [do they??]). Certainly no one ought to expect a semantic theory to compute this without even a statement of what the context is. And if we *are* supplied with a statement of what the context is, so that we *have* a way to characterize how it affects the literal meanings of words and the literal semantic effect of syntactic operations so as to produce the occasion meanings of these words and operations, then we might indeed find ourselves in a position to give a compositional semantic account of occasion meaning from the literal meaning. It is that possibility which we will now explore.

As it affects the examples of this section, the idea is this: the theory of context takes as input the literal meanings of the atoms of the language (for practical purposes, let us just say the words and not worry about morphosyntax) and the literal semantic effects of the various syntactic constructions, and for each specific context (that is, each updated model that is generated in the sentence-by-sentence processing of a discourse, in the picture developed earlier) yields as output what is the occasion meaning of that word or construction in that model. (It is certainly possible that an update to a model will not effect a change to some or most words/constructions, and they will retain their occasion meaning from preceding models; and this may extend all the way to the first, so-called neutral model, which comprises just the literal meanings of the words/constructions.) The compositional machinery uses the occasion meanings thus supplied to compute the occasion meaning of the sentence under consideration. In the case of (10.8) we have:

(10.10)   A previous model has identified the referent of *that* as BD, and has settled that the meaning of *prospect* is BUSINESS-DEAL. The computational machinery then computes: NEW-FOR-A-BUSINESS-DEAL (BD) ∧ BIG-FOR-A-BUSINESS-DEAL (BD)

This is the strategy pursued in (Recanati 2012: 188–90), and to some extent in this present chapter.[12] (Differences will emerge later). Recanati postulates that there is a function which will take a context and return a modulation function, *mod*, that in turn will use an interpretation function to generate the occasion meanings of each term in the sentence as spoken in that context. And since a complex expression might itself be a part of a more encompassing complex expression, both of which need to have this mechanism operate on them, Recanati wants to employ some recursive mechanism that uses the modulations and interpretations of sub-expressions to generate them for the encompassing ones. He refers approvingly to the method of Pagin and Pelletier (2007: esp 46–50) for this generalization. Further discussion of Recanati's notions of context and compositionality follows below in the next Section.

10.4. More on contexts

10.4.1. *Recanati on contextual compositionality*

Recanati (2004) distinguished two types of contextual processes: *Saturation* and *Modulation*. Saturation is the process by which indexicals and any 'free variables in logical form' are assigned a contextual value; Modulation is the set of processes whereby the meaning of an expression is changed into some other meaning by a 'pragmatic function'. There are sub-types of modulation, of which Recanati talks

---

[12] Recanati refers to Sag (1981) as initiating the formalization of this sort of approach.

especially about *metaphoric interpretation, metonymic interpretation*, and *free enrichment*. (Metaphor and metonymy are fairly standard operations in the rhetoric literature; by 'free enrichment' Recanati means methods which map the meaning of an expression onto some more specific meaning.)

While 'everyone' is happy to allow for Saturation by context (to give content to indexicals), Recanati wants to allow Modulation also to work freely to give occasion meaning to *any* expression at all. And in his (2012), Recanati raises the question of whether this can be done in a compositional manner.

Of course, it is no longer the standing meaning that we are interested in, and which compositionality was initially developed to describe. It is instead the occasion meaning. Recanati's (2012) view is that we can allow semantic compositionality into pragmatics by this ploy: '…the contextualist [holds that] we should do our best to account for the intuitive truth- and satisfaction-conditions of utterances, and to that effect we may have to liberalize the notion of meaning/content to the point of blurring the semantics/pragmatics distinction.' (2012: 188).[13]

In Recanati's view, the occasion meaning of a whole is a function of the *modulated* meanings of its parts and the manner in which they are combined, much as I had suggested above where I had 'occasion' in place of 'modulated'. To compute occasion meaning, Recanati introduces 'a function *mod*, taking as argument an expression *e* and the context c in which it occurs: the value of *mod* is the particular modulation function g that is contextually salient/relevant/appropriate for the interpretation of that expression in that context.' (2012: 188)

This is added to the 'standing interpretation function', $[\![\phi]\!]_c$ that yields ϕ's standing meaning in context c (the content of ϕ after saturation of the indexical elements of ϕ). So: the relevant occasion-meaning-in-context-c of a lexical item e is described as:

(10.11) $mod(e, c)([\![e]\!]_c)$

That is, the *mod*-function associated with e in c applied to the content of e in c (which in turn is the saturation of e by c).

For the syntactic complexity rules we get:

(10.12) $[\![\phi * \psi]\!]_c = f(mod(\phi, c^1)([\![\phi]\!]_{c^1}), mod(\psi, c^2)([\![\psi]\!]_{c^2}))$

where c^1 and c^2 are 'the relevant sub-contexts' of context c. If we specify that $mod(\phi, c^1)$ is g_1 and $mod(\psi, c^2)$ is g_2, then this becomes:

(10.13) $f(g_1([\![\phi]\!]_{c^1}), g_2([\![\psi]\!]_{c^2}))$

---

[13] This seems to follow the various positions that identify meaning with use, such as (Quine 1987: 211): 'The separation between semantics and pragmatics is a pernicious error.'

(Recanati (2012: 183 n 4) allows that some syntactic rules might also be modulated, but as I said before, he doesn't follow this up.) Literal/standing meaning—which occurs in a 'neutral' context—is the identity *mod*-function.

Note that Recanati thinks that *mod* is itself functional: the *mod* of larger syntactic units within an utterance is a function of the *mod*s of the simplest parts, the words. So, the occasion meaning of *The big prospect* is some function of the occasion meanings of *The*, *big*, and *prospect*. There is no new information added by context that would further influence the occasion meaning of the entire NP.

10.4.2. *What should we say about Recanati's picture?*

One thing is: *mod* is pretty magical!!

It takes *any* word (also, at least in theory, it takes *any* syntactic operation) in *any* context and returns a function that tells you how to alter the standing meaning of that word (or operation) in that context to give the occasion meaning (in that context). Not only is it super powerful that way, but also: it's a *function*—*Any* time you give it the same word in the same context, it will feed back the very same function.

To do this involves individuating contexts *very* finely however, since anything that might *ever* be relevant to altering the standing meaning of some context would have to be used to individuate *all* contexts. Winks of the eye, for example. Smiles rather than frowns. Being a habitual jokester. Kicks under the table. Being from Canada.

So, I think we cannot really maintain Recanati's picture. It's tempting to say that the 'magical nature' of *mod* detracts from, or maybe even eliminates, the entire reason that one wants to maintain compositionality in the first place – the Arguments from the Infinity of Language. Recanati's domain of contexts (which is the basis of *mod*) simply is not finite in the way that the lexicon and syntactic-semantic rules are finite in number.[14, 15]

10.4.3. *A slightly different picture*

In the past (e.g., Pelletier 1994) I have argued that compositionality is false in the classic form in which it is stated. But I'm not going to rehearse that now. Instead, I merely want to present the different picture (which I call *semantic groundedness*) that I prefer to semantic compositionality, and argue that some version of it ought to replace Renanati's attempt to use compositionality for pragmatic or occasion meaning.

---

[14] Recanati (2012: 190–1) claims that the domain of contexts *is* 'always finite'. Perhaps this is true for any specific context, but his general mechanism would seem to require that the entire realm of all domains should be finite. And this seems questionable to me. More discussion about this point is further in footnote 16.

[15] In discussing this issue in his review of (Recanati 2010)—where the same claims are made—Bach (2011: 7) judges that Recanati's pragmatic modulation 'does not suggest anything in the way of a procedure for determining what these functions are in any particular case. . . . It is hard to see how this isn't tantamount to giving up even a weak form of compositionality.'

As a starting point, let us note that compositionality is closely related to recursive definitions:

$$\phi(n) =_{df} a, b, c \ldots, \qquad \text{if } n = 0, 1, 2, \ldots \quad \text{[basis clause]}$$
$$=_{df} \psi(\phi(n-1)) \qquad \text{otherwise} \qquad \text{[inductive clause]}$$

For example, the usual definition of the factorial function follows this characterization:

$$n! =_{df} 1, \qquad \text{if } n = 0$$
$$=_{df} n*(n-1)! \quad \text{otherwise}$$

Similarly with the case of 'complete induction', where we have:

$$\phi(n) =_{df} a, b, c \ldots, \qquad \text{if } n = 0, 1, 2, \ldots \qquad \text{[basis clause]}$$
$$= \psi(\phi(x), \phi(y), \ldots), \text{otherwise}, (x, y, \ldots < n) \quad \text{[inductiveclause]}$$

For example, the usual definition of a formula follows this characterization (I use the inductive clause only for $\rightarrow$):

ϕ is a formula,   if ϕ is atomic
$(\phi \rightarrow \psi)$ is a formula,   if ϕ and ψ are formulae

When people think of compositionality they tend to think of a complicated version of inductive recursion, in either its composition mode or its decomposition mode.

Inductive recursion:
[**composition mode**] Start with some base case and continue to apply the rule/ function to newly-created values.

[**decomposition mode**] Start with some complex item and continuously decompose it until you reach the basis case.

Note that in applying this recursive function in decomposition mode to any complex item, determining its value involves determining values of items strictly smaller. And that is the fundamental intuition behind semantic compositionality.

10.4.4. *Semantic groundedness*

But what's wrong with doing it like this instead?

Call a function *General Recursive* (as opposed to inductively recursive) if every evaluation of an argument that the function is applied to eventually or ultimately reaches a basis case. Note that in this characterization it is *not* the case that each argument of the function is evaluated by applying the function to 'smaller' values of the argument. The value of some function f need not apply to the immediately smaller value, nor indeed to any smaller value. For example, $f(n)$ might require the evaluation of $f(m)$, where $m = 2n$, and in turn, this itself requires the evaluation of $f(j)$,

where $j = m/8$, and so on. The only requirement is that the evaluation always terminate in some basis case.

For example, consider two ways of defining a truth assignment, f^*, to sentences of propositional logic. In both ways we are given an assignment of truth values to the atomic sentences, f, and we extend this. Extending it in an inductive recursive ('compositional') way goes as follows:

$$
\begin{aligned}
f^*(p) &= f(p) & \text{if } p \text{ is atomic} \\
f^*(\neg\,\phi) &= 1 & \text{if } f^*(\phi) = 0 \\
&= 0 & \text{if } f^*(\phi) = 1 \\
f^*((\phi \rightarrow \psi)) &= 1 & \text{if } f^*(\phi) = 0 \text{ or } f^*(\psi) = 1 \\
&= 0 & \text{if both } f^*(\phi) = 1 \text{ and } f^*(\psi) = 0 \\
f^*((\phi \wedge \psi)) &= 1 & \text{if } f^*(\phi) = 1 \text{ and } f^*(\psi) = 1 \\
&= 0 & \text{if either } f^*(\phi) = 0 \text{ or } f^*(\psi) = 0 \\
f^*((\phi \leftrightarrow \psi)) &= 1 & \text{if } f^*(\phi) = f^*(\psi) \\
&= 0 & \text{if } f^*(\phi) \neq f^*(\psi)
\end{aligned}
$$

But this is not the only way to define f^*. Here's a general recursive ('non-compositional') way:

$$
\begin{aligned}
f^*(p) &= f(p) & \text{if } p \text{ is atomic} \\
f^*(\neg\,\phi) &= 1 & \text{if } f^*(\phi) = 0 \\
&= 0 & \text{if } f^*(\phi) = 1 \\
f^*((\phi \rightarrow \psi)) &= 1 & \text{if } f^*(\phi) = 0 \text{ or } f^*(\psi) = 1 \\
&= 0 & \text{if both } f^*(\phi) = 1 \text{ and } f^*(\psi) = 0 \\
f^*((\phi \wedge \psi)) &= 1 & \text{if } f^*(\phi) = 1 \text{ and } f^*(\psi) = 1 \\
&= 0 & \text{if either } f^*(\phi) = 0 \text{ or } f^*(\psi) = 0 \\
f^*((\phi \leftrightarrow \psi)) &= f^*((\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi))
\end{aligned}
$$

Even though the way f^* is defined for the $(\phi \leftrightarrow \psi)$ case does not appeal to the subparts of that formula, we *know* that it will always terminate. So, as defined here, f^* is a general recursive function.

Note that the inductively recursive functions are a subset of the general recursive ones. Note also that the general recursive functions that are not inductively recursive are also not compositional in any straightforward sense.

A slightly different group of functions are the ones I call *semantically grounded*. These functions are like the general recursive ones in that they allow function evaluation to use values that are not strictly smaller (or 'part of') the argument being evaluated. However, they are not general recursive in the literal sense because they might not always terminate in a basis case for every *possible* argument that the function can be applied to. Being a semantically grounded function is more of an empirical or practical notion, where the function always terminates in a

basis case *in all the cases we have encountered* or *in all the cases of practical interest*.[16] The non-compositional $f*$ is semantically grounded: we know that it will always terminate (for whatever reason we know this), and so it does terminate in all the cases of interest. Other semantically grounded functions perhaps don't always terminate; but in all the cases we are aware of, or in all the cases we are interested in, they *do* terminate.

Looking just at language (even though I think the point also holds of other areas, such as metaphysics), recall that 'atomistic theories' eventually find some 'grounding instances', as opposed to those contextualist theories that appeal to a context that is itself a 'whole' in which atomic parts cannot be discerned. (As the holists would say, 'Every part of the context depends on all the other parts in order to have the force on meaning that it has'.) The semantic groundedness theory falls into the atomistic camp: although one might need to look to things other than a (syntactic) object's parts, eventually these each lead to a basis case—a fundamental grounding point. In the case of evaluating '$\phi \leftrightarrow \psi$' we looked to a conjunctive formula, which involved looking at conditionals. But these conditionals were directly 'grounded'. So this was a case of an atomistic, but non-compositional theory. It is a case of semantic groundedness.

10.4.5. *Semantic groundedness and the infinity of language*

Semantic compositionality presents the picture that there are a finite number of atomic items of a language (and their meanings) and there are a finite number of syntactic ways to put items together (each with its meaning-effect). If the syntax allows an unbounded number of things to be constructed from the atomic items it is because the syntactic rules allow recursive re-application. But then we could *learn* this, since all we need learn are the finite number of atomic items and syntactic rules; and we can compute the meanings of any of these infinite number of things (answering the *understandability* criterion).

But surely semantic groundedness gives much the same picture: it is an *atomistic theory*, and so it claims to eventually find some 'grounding instances', as opposed to contextualist theories. Although one might need to look to things other than a (syntactic) object's parts to compute or construct meanings of a complex item, eventually these each lead to a base case—a fundamental grounding point. Or at least, it does lead there in all the cases of interest, or all the cases we've encountered.

---

[16] Recanati (2012) seems also to like this 'empirical' or 'practical' interpretation of how context and compositionality can interact. He says:

> Contextual modulation provides for *potentially* unending meaning variation, but never gives rise to any *actual* unending meaning variation. Meaning eventually stabilizes, making compositionality possible, because the (linguistic as well as extralinguistic) context, however big, is always finite. (190–1) (Emphasis in original)

And while his picture of 'modulation plus compositionality' is somewhat different from my semantic groundedness, on this more general point we seem in agreement.

The picture I have for semantic groundedness allows that a semantic evaluation could bring into play all kinds of facts, information, context, inferences, world knowledge, etc, even when they are not 'meanings of the parts or of their modes of combination' (so long as it is always grounded—or maybe, grounded in cases where actually employed). If an evaluation is grounded, then this provides as good an explanation of the learnability/understandability facts as compositionality does; indeed, it's essentially the same explanation. But given that the application is to features of context, including the long list of related factors that I just mentioned, perhaps it should be called *pragmatic* groundedness.

The idea is that there is a bunch of independent factors that tap many different information sources, and not just one specific source—as Recanati's use of a function on contexts that then generates specific *mod*-functions is. Some sources are relevant in some contexts but not in others, and there may be no way to tell in advance which ones will be relevant in any particular context. In the picture proposed by pragmatic groundedness, all that is required is that any of these sources be ultimately grounded in *some* sort of base. In my view, as in Recanati's, these results then feed into the compositional framework, more or less along the lines that Recanati advocates. But: they are *not* the result of applying some inductively recursive mechanism to a finitely-based domain, as Recanati claims, and are thus not compositional overall.

So as a consequence, I am recommending that Recanati not think of his *mod* as a function. Instead, he should view it as an information source. The requirement on the source is that it be grounded, or grounded in the cases that are currently under consideration in any specific speech situation—that is, that it not make use of information sources that in turn somehow involve the output of that original source.

This is not a very big difference so far as the technical development goes. But it is a big alteration in the overall viewpoint. Although the theory remains atomistic—as I think is Recanati's main goal—it is not compositional.

10.5. What does this say about brevity?

Although one could write a long disquisition on the topic of this section, I limit myself to pointing to just one example where the outlook suggested above has important consequences for a theoretical account of why we can be brief-in-context. I think the considerations mentioned can easily be applied to other cases, and thus supports the general conclusion at the end of this section.

The main thrust of this chapter, ignoring certain ancillary issues that were discussed, has been the following: (a) it is clear that features present in context can somehow allow for brevity of conversation, and we would like a theoretical account of that; (b) however, this must be tempered by the fact that any such theoretical account must honour the facts concerning the Infinity of Language (which rules out a variety of accounts that go under the title of 'holistic') in their explanation of how to

describe the contextual meaning of any piece of language; (c) in turn, composition-ality—the most straightforward way to honour the Infinity of Language—places restrictions on just how brief one can be even in a very rich context, and it is also difficult to integrate compositionality with the richness of the possible features of context that can give rise to brevity. So, (d) I described an atomistic alternative to compositionality called *pragmatic grounding*. This alternative inherits composition-ality's explanation of the Infinity of Language, at least for the most notable or relevant cases, and furthermore opens up for exploitation many more features of a context that can be used in generating briefer remarks than compositionality would allow.

To see how this plays out, consider the well-worn example (inspired by Nunberg (1979)), 'The ham sandwich did an eat-and-dash', perhaps followed by 'No, no—that was the deep-fried pickle', said in a restaurant among members of the staff who have previously experienced conversations that are similar in one way or another to the present scene. (Maybe once some customer was burned when touching an extremely hot plate and the manager demanded 'Who served the plaid jacket a too-hot plate!?'; perhaps the waitperson replied 'No, it wasn't me: the plaid jacket was the margarita pizza, and had no problems'; or they discussed how 'The police no longer respond to gas-and-dash incidents', etc.)

As I pointed out in Section 10.2.3, to make it plausible that the interlocutors can generate such utterances with knowledge that they correctly describe what they want to assert (and that the audience will understand them), we somehow need to pass from the literal or standing meaning of the terms to the occasion meaning, or conversely. The occasion meaning of 'the ham sandwich' in this context is (something like) the literal meaning of 'the person in this restaurant who ordered and ate a ham sandwich', and the occasion meaning of 'did an eat-and-dash' is (something like) the literal meaning of 'ran out of the restaurant without paying for his food'. If all the speakers involved could know that the occasion and literal meanings of the terms were as just stated, then they could either: (a) generate the literal meaning of the sentences from the occasion meanings of the various terms (which we are hypothesizing that they know) and (b) apply compositionality on the literal meanings to generate an understanding; or else they could (c) stick with the occasion meanings of the terms and apply compositionality on the occasion meanings to generate the relevant occasion meaning understanding. The problem is to find a mechanism that would allow us either to correlate the literal and occasion meanings (so as to do (a) and (b)) or else to generate the occasion meanings from the literal meanings plus contextual features (so as to do (c)).

The strategy in both Recanati's (2012) proposal and mine is to generate the occasion meanings from the literal meanings plus context so as to do (c). The differences are embedded in how we propose to carry this out. But if either proposal can be successfully developed, then there would be a genuine theoretical account of how brevity-in-context is possible. Without some such account we would be open to the sorts of objections that Swift raised to the Professors of Language in the Grand Academy of Lagado.