

Context, Content, and Compositionality **François Recanati**

Extended synopsis

State of the art, objectives and methodology of the project — Over the past fifteen years, I have argued that the effects of context on content go well beyond what is standardly acknowledged in semantics. This view is sometimes referred to as ‘Contextualism’ or (more technically) ‘Truth-Conditional Pragmatics’ (TCP). The key idea is that the effects of context on content need not be traceable to the linguistic material in the uttered sentence. Some effects are due to the linguistic material (e.g. to context-sensitive words or morphemes which trigger the search for contextual values), but others result from ‘top-down’ or ‘free’ pragmatic processes that take place not because the linguistic material demands it, but because the literal meaning of the sentence requires adjustment or elaboration (‘modulation’) in order to determine a contextually admissible content for the speaker’s utterance. Such processes, e.g. free enrichment, sense extension, metonymic transfer, etc., arguably affect the intuitive truth-conditions of utterances even though they take place for pragmatic reasons, without being triggered by the linguistic material in an obligatory manner.

The claim which TCP makes regarding the role of free pragmatic processes in the determination of intuitive truth-conditional content is an empirical conjecture about natural language. Other philosophers of language (Stanley 2000, 2007, Szabo 2000) have made the opposite conjecture, more in line with traditional ways of thinking about meaning and truth-conditions. Since it gives up those traditional assumptions, TCP sounds revolutionary, and many theorists (e.g. Predelli 2005) are suspicious of it because they take it to threaten the very enterprise of semantics. In the literature, one often finds arguments to the effect that, if Contextualism is right, then systematic semantics becomes impossible. More precisely, the claim that is often made is that TCP is incompatible with the Principle of Compositionality, upon which any systematic semantics must be based. Thus Jerry Fodor argues as follows. Even though the sense of an equivocal expression depends upon the linguistic context (e.g. the sentence) in which it occurs, this does not prevent it from serving as building block in the construction of the meaning of that sentence, as compositionality requires. The expression *has* a definite sense, which the context reveals, and that sense has the stability required to serve as building block. An insuperable problem arises, however, as soon as we generalize this type of context-dependence, as TCP does. Take an expression *a* and assume its modulated meaning depends (inter alia) upon the meaning of its complex host. The meaning of the host itself is liable to vary, because of modulation, depending on *its* host : and so on indefinitely. So the meaning of *a* will never stabilize : there is unending equivocation, ‘equivocation that can’t be resolved’, and it ‘undermine[s] the compositionality of English’ (Fodor 2003 : 99).

The aim of this project is to defend Contextualism/TCP by demonstrating that it is *not* incompatible with the project of constructing a systematic, compositional semantics for natural language. This demonstration is of paramount importance given the current predicament in the philosophy of language. We are, as it were, caught in a dilemma : formal semanticists provide compelling arguments that natural language must be compositional, but

contextualists offer no less compelling arguments to the effect that « sense modulation is essential to speech, because we use a (mor or less) fixed stock of lexemes to talk about an indefinite variety of things, situations, and experiences » (Recanati 2004 : 131). What are we to do, if modulation is incompatible with compositionality ? Our aim is to show that it is not, and thereby to dissolve the alleged dilemma.

As far as methodology is concerned, we will proceed, as philosophers typically do, by *outlining possible accounts* that satisfy the compositionality requirement while remaining faithful to the spirit of Contextualism. In the second half of the project (workpackages 3 and 4) we will show how free pragmatic processes can be understood within a computational theory of mind à la Fodor, and we will do that, once again, by *outlining possible theories*, thereby opening the way to empirical investigation of fundamental issues relating to the language/thought interface.

Background : the compositional procedure – The compositionality idea is the idea that semantic interpretation proceeds in two steps. Simple expressions are interpreted by means of *lexical rules*, which assign meanings to them directly. Complex expressions are interpreted by means of *compositional rules*, which assign meanings to them indirectly, as a function of the meanings of their parts.

For any simple expression a , the associated lexical rule says that the interpretation of a is a certain entity m :

$$(1) \quad I(a) = m$$

There will be as many rules of this sort as there are simple expressions (or, rather, readings of simple expressions) in the language. Since the number of simple expressions and the number of readings which an ambiguous expression has are both finite, it is, in principle, possible for a finite mind to get to know the meanings of all simple expressions of the language by learning each of the lexical rules that are associated with them in this way.

The syntax of natural language is such that (because of recursivity) the number of complex expressions is *not* finite : for any expression of whatever complexity is it always possible to construct a more complex expression. So it would not be possible for a finite mind to get to know the meaning of all expressions of the language, simple *or* complex, by learning that meaning directly. If we only had rules like (1) to interpret a linguistic expression, there would have to be an infinite number of them, and we could not learn them. So we need a different type of rule than (1) for interpreting complex expressions.

Just as the number of simple expressions is finite, the number of ways in which distinct expressions can be put together so as to yield a complex expression of the language is finite. In other words, there is a finite number of syntactic rules, through which an infinite number of complex expressions can be generated. The solution, then, is to pair each syntactic rule with a semantic rule of a new sort – a compositional rule. A compositional rule is something like

$$(2) \quad I(a*b) = f(I(a), I(b))$$

where '*' stands for an arbitrary mode of combination. The rule says that the interpretation of the complex expression $a*b$ is the value of a certain function f when it takes as arguments the interpretation of a and the interpretation of b .

Thanks to rules of this sort, it is possible to compute the meaning of an expression of whatever degree of complexity on the basis of the meanings of its parts. If the parts are simple, their meanings $I(a)$ and $I(b)$ will be given directly by lexical rules such as (1). If the parts are themselves complex, their meanings will themselves be derivable via compositional rules such as (2).

In this framework, the meaning of a complex expression only depends upon two things : the meanings of its immediate constituents (the simpler expressions into which it can be analysed), and the way they are put together. Nothing else counts. In particular, the meaning of an expression does not depend upon the meanings of other expressions that are not its constituents, even if they occur in the same sentence or discourse. Nor can the meaning of a given expression depend upon the meaning of a more complex expression in which *it* occurs as a constituent.

Workpackage 1 : Accounting for semantic flexibility within a compositional framework — As we have just seen, 'top-down' or 'lateral' influences on meaning are ruled out by the compositional procedure. Yet such influences are precisely what we find, and this raises a problem which our first aim is to solve.

Let us say that a language exhibits *semantic flexibility* if the following condition is satisfied : in that language, the meaning of a word may vary from occurrence to occurrence, and it may vary, in particular, as a function of the other words it combines with. Through semantic flexibility, the meaning of an expression may well depend upon the meaning of the complex in which it occurs (top down influence), and it may also depend upon the meaning of the other words that occur in the same complex (lateral influence). One of the authors who have insisted that natural languages exhibit semantic flexibility is Jonathan Cohen in a series of papers in which he criticizes mainstream approaches in semantics. He gives examples like the following :

Consider 'drop'... in the sentences

(3) Most students here drop geography in their final year

(where 'drop' means 'drop studying'),

(4) Most students here drop geography lectures in their final year

(where 'drop' means 'drop attending'),

(5) Most students here drop geography lectures reading assignments in their final year

(where 'drop' means 'drop executing')

(6) Most students here drop geography lectures reading assignments library-fees in their final year (where 'drop' means 'drop paying'),

and so on indefinitely. If we accept that a sentence can be as long as we please, then there 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. (Cohen 1986 : 227-8)

According to Cohen, the verb 'drop' takes on a different meaning in each of (3) to (6), and one of the things that determine the meaning it takes on is the noun phrase it combines with. A similar type of example is provided by John Searle, a forerunner of Contextualism :

The sort of thing that constitutes cutting the grass is quite different from, e.g., the sort of thing that constitutes cutting a cake. One way to see this

is to imagine what constitutes obeying the order to cut something. If someone tells me to cut the grass and I rush out and stab it with a knife, or if I am ordered to cut the cake and I run over it with a lawnmower, in each case I will have failed to obey the order. (Searle 1980 : 222-223)

According to Searle, 'cut' means something different — has different satisfaction conditions — in 'cut the grass' and in 'cut the cake' ; and that is because the meaning which the verb 'cut' takes on on a particular occurrence depends, inter alia, upon what is said to be cut.

The examples I have given involve a transitive verb the (exact) meaning of which depends upon the noun phrase that serves as its complement. An even more productive class of examples involves adjectives the (exact) meaning of which depends upon the noun they modify. A good car is not good in exactly the same sense in which a good house is ; a piece of luggage is not light in exactly the same sense in which a sound is light ; a big mouse's way of being big differs, to some extent, from the way in which a big elephant is big ; a pink grapefruit is not pink in the same way — under the same aspect — as a pink raincoat ; and so on and so forth. In all cases the basic meaning of the adjective is fleshed out differently according to the noun it modifies.

Semantic flexibility and compositionality, as I have characterized them, seem to be mutually exclusive properties. As Jerry Fodor puts it,

The compositionality thesis says that complex representations inherit their content from simple ones, *not vice versa*. But the contextualist thesis says that the content of a simple [representation] depends (inter alia ?) on which complex [representation] it's embedded in. Clearly, it can't be that both are true. Something's gotta give. (Fodor 2003 : 96-7)

Cohen seems to agree. According to him (1986 : 230), because of semantic flexibility, « we cannot construct a semantics for any natural language along the same lines as a semantics for a formal system of any currently familiar kind. Projects like Davidson's or Montague's cannot succeed ».

In the first stage of the project, I will *question the assumption that semantic flexibility is incompatible with compositionality*. I will proceed by outlining, and comparing, two ways of accounting for semantic flexibility within a compositional framework. The most promising account treats semantic flexibility as ultimately a matter of context-sensitivity. On this view, to be fleshed out in detail, the content of the complex phrase is a function of the contents of its parts, in a strictly bottom-up manner ; but the content of the parts is, or may be, context-dependent, and the linguistic context in which an expression occurs is an aspect of the context which may influence its content. Lateral and top-down influences are therefore possible — the content carried by a particular expression may depend upon the other expressions with which it combines — but this is compatible with the fact that the content of the whole depends upon the contents of its parts in a strictly bottom-up manner.

Workpackage 2 : Accounting for modulation within a compositional framework — In *Literal Meaning* and elsewhere I drew a systematic distinction between two types of contextual process possibly affecting truth-conditions: the (mandatory) process of 'saturation' through which indexicals and free variables in logical form are assigned a contextual value, and the (optional) process of 'modulation' through which the meaning *m* of an expression is mapped to a

distinct meaning $g(m)$, where 'g' is a pragmatic function. Metaphorical and metonymical interpretations result from the operation of such pragmatic functions, and the argument to the function may be the meaning of *any* expression, whether or not it is 'context-sensitive' in the standard sense in which indexicals and semantically under-specified expressions are. Another type of pragmatic function, involved in so-called 'free enrichment', maps the meaning of an expression to a more specific meaning.

Now, if semantic flexibility is to be accounted for by appealing to contextual processes, there is no reason to restrict the type of contextual process at issue to saturation. We may get the sort of contextual influence on the interpretation of a lexical item which gives rise to the phenomenon of semantic flexibility even if the expression whose interpretation contextually varies in this way is not indexical or context-sensitive in the standard sense. Just as the value contextually assigned to an indexical or free variable may be influenced by the linguistic context, the modulated value which a given expression takes in context may also be influenced by the linguistic environment, that is, by the other words with which the expression combines. Consider 'The city is asleep', an example I discussed in *Literal Meaning* (pp. 34-36). Because of the apparent category violation (a city is not the sort of thing that sleeps) either 'asleep' must be interpreted in a metaphorical or extended sense as meaning QUIET AND SHOWING LITTLE ACTIVITY, or 'the city' has to be interpreted metonymically as referring to the inhabitants of the city. Either way, how we interpret one expression depends upon how we interpret the other. This is semantic flexibility once again, but of course we do not want to account for that type of example in terms of indexicality or context-sensitivity, but in terms of modulation.

In the second stage of the project, we will *explore the ways in which we can make room for modulation within a broadly compositional approach*. The basic idea is that we should allow for compositional rules of the following sort :

$$(2^*) \quad I(a*b) = f(g_1(I(a)), g_2(I(b)))$$

In that formula the 'g's are free higher-order variables ranging over available pragmatic functions (including identity, which gives us the 'literal' case). The formula says that the semantic value of a complex phrase a*b is a function of the pragmatic values of the parts, where the 'pragmatic values' in question are what we get when we subject the literal semantic values of the parts to pragmatic modulation. Thus 'lion', in the right context, comes to mean 'representation of lion', 'straight' comes to mean 'approximating straightness', and so on and so forth. Those 'modulated' meanings are the building blocks out of which the meaning of complex phrases like 'stone lion' or 'pretty straight' are built. (A stone lion is not a real lion, and something that is pretty straight is not really straight.)

Workpackage 3: Modulation as a 'syntactic' process — In the third stage of the project we will *explore the implications of this sort of view for the way we think of the language/thought interface*, and in particular, for the notion, or notions, of 'logical form' used in current linguistic theory. The logical form of a sentence is taken to be a (covert) level of syntactic structure, but the logical properties relevant to semantic interpretation (quantifier scope, anaphoric dependencies, etc.) are supposed to be fixed at that level. Indeed, some authors, like Chierchia (1999), take 'logical forms' to be *conceptual* representations — representations in the 'language of thought' — yet conceptual representations

that are strictly determined by the grammar. On this view logical forms belong both to the linguistic system and to the conceptual system, and serve as interface between the two systems. As conceptual representations they can be elaborated or modified through non-linguistic considerations, and this is where free pragmatic processes possibly come into the picture. So we will elaborate what I call the 'syntactic' construal of free pragmatic processes, according to which they map conceptual representations directly determined by the grammar to further conceptual representations (which I call 'modified logical forms'). The main advantage of this construal is that it enables us to raise a fundamental issue that has never received the systematic treatment it deserves.

Like Chierchia, relevance theorists view logical forms as the interface between language and thought : they are (partial) conceptual representations that are determined solely by the language system. But these logical forms are not endowed with the determinateness of full-fledged conceptual representations. They lack important logical properties that will only be determined at a further level of conceptual elaboration and, for that reason, they cannot be given a truth-theoretic interpretation. Fodor holds the same view : he insists that what gets compositionally interpreted by means of recursive truth-theoretic procedures is not what is determined strictly by the grammar but the modified logical form which is a syntactic representation in the language of thought, and which is affected by pragmatic processes and world knowledge (Fodor 2001 : 12-13). Even within the mainstream generative tradition, there are authors who think that certain logical properties relevant to semantic interpretation are not fixed by the syntax at LF and must therefore be dealt with at the 'conceptual' level. Thus the LF which May 1985 ascribes to a sentence like 'What did everyone bring ?' is ambiguous with either quantifier capable of taking wide scope. Jackendoff goes even further. For him, the language system and the conceptual system do not intersect in the way suggested by the mainstream 'logical form' idea. Rather, the two systems are disjointed, and additional 'rules of correspondence' are needed to bridge the gap between the syntactic structures of language and syntactic representations in the language of thought (Jackendoff 1993 : 31). So there is no reason why we should expect any level of linguistic representation to display the logical properties (like quantifier scope) which characterize conceptual representations (Jackendoff 2002 : 270).

Does the linguistic system output a representation which can serve directly as input to truth-theoretic semantics (as well as serving as input to the pragmatic processor)? Or, is only the output of the pragmatic processor semantically interpretable ? One of our tasks in the third stage of the project will be to determine what sort of empirical evidence can be brought to bear on that issue.

Workpackage 4 : The 'syntactic' construal without the language/thought duality — The assumption that there are two distinct systems (the language system and the conceptual system) can be questioned. According to some philosophers, there is a *single* system (e.g. Carruthers 1996, Ludlow 1999 : 164-9) : thought is nothing but 'inner speech' ; or at least, it is underpinned by the linguistic system, which provides the structures for the articulation of thoughts. Can we still appeal to the notion of modified logical form, if we do not invoke a second system in addition to the language system ? That is the question we will deal with in the last phase of the project. We will provide an affirmative answer, based on re-interpreting a rather familiar idea.

According to many authors, whenever a sentence is uttered and assigned,

in context, a 'modified logical form' resulting from the operation of free pragmatic processes, the modified logical form in question is the logical form of *another sentence*, that which runs through the speaker's and/or the interpreter's mind. There are two versions of that view in the literature. One of them can be found in the writings of theorists such as J. Katz and K. Bach : when a sentence is used to convey something different from - e.g. more determinate than - its literal interpretation, pragmatic processing maps that sentence to some other sentence that was not uttered but might have been. The other possible interpretation of the view is less familiar, but it has recently found advocates in the linguistic and philosophical community (Marti 2006, Stern 2006). It goes like this. In the relevant examples, there are, indeed, two sentences s_1 and s_2 , where s_2 corresponds to the actual interpretation of the utterance ; but it is a mistake to think that s_1 is uttered, while s_2 is only mentally tokened. What is uttered actually is... s_2 ! On this view the two sentences s_1 and s_2 are *phonetically undistinguishable*, because what differentiates them are only *covert* elements in the logical form of s_2 , which are missing in the logical form of s_1 . These covert elements manifest themselves in the semantic interpretation and reveal that the uttered sentence is s_2 , not s_1 . On this picture, which we will elaborate and whose consequences we will explore, the effect of so-called free pragmatic processes are nothing but the effects of semantically interpreting covert elements in logical form. So, in a sense, there is no free pragmatic process. In another sense, however, there are such processes, but they must be redescribed and accounted for in terms of the free generation of pragmatic variables in the syntax. (Elaborating this approach may prove useful even if one sticks to the standard view, according to which the language system and the conceptual system are distinct. For this may help us discover criteria for discriminating between free pragmatic processes which have syntactic reality within the language system, and others whose putative 'syntactic' reality lays entirely outside the language system, in the conceptual system.)