



HAL
open science

Literalism and Contextualism: Some Varieties

François Recanati

► **To cite this version:**

François Recanati. Literalism and Contextualism: Some Varieties. Gerhard Preyer. Contextualism, Oxford University Press, 2003. ijn_00000371

HAL Id: ijn_00000371

https://hal.science/ijn_00000371

Submitted on 8 Sep 2003

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Literalism and Contextualism: Some Varieties

FRANÇOIS RECANATI

According to the dominant position in the philosophy of language, we may legitimately ascribe truth-conditional content to sentences, independently of the speech act which the sentence is used to perform. This position, which I call 'Literalism', contrasts with another view, reminiscent of that held by ordinary language philosophers half a century ago. That other view, which I call 'Contextualism', holds that *speech acts* are the primary bearers of content. Only in the context of a speech act does a sentence express a determinate content.¹

Both Literalism and Contextualism come in many varieties. There are radical, and less radical, versions of both Literalism and Contextualism. Some intermediate positions are mixtures of Literalism and Contextualism. In this paper I will describe several literalist positions, several contextualist positions, and a couple of intermediate positions. My aim is to convince the reader that the Literalism/Contextualism controversy is far from being settled.

In the first section, I will look at the historical development of Literalism. We will see that this development reveals a gradual weakening. The question that naturally arises is: How far can we go in this direction? Where will this tendency ultimately lead us? And the obvious answer is: to Contextualism. In the second section I will describe the steps which, from a critique of the currently dominant literalist position (Minimalism), can lead to Contextualism. In the last three sections I will describe various contextualist positions, and I will discuss possible literalist replies to the contextualist challenge.

I. The Development of Literalism

Indexicality raises a *prima facie* difficulty for Literalism — a difficulty that was emphasized by its contextualist opponents. Indexical sentences possess a determinate (truth-evaluable) content only when uttered. Hence it is not obvious that such sentences, *qua* grammatical entities, possess content. As the ordinary language philosophers used to insist, we must draw a distinction between the sentence and the

¹ As James Conant (1998) pointed out, this is a Wittgensteinian extension of Frege's Context Principle.

statement it is used to make. The content is the content of the statement, and only derivatively that of the sentence that is used to make that statement.

But indexicality is a feature of natural language sentences, and the philosophers in the literalist tradition were not originally concerned with natural language. They were primarily concerned with the formal languages of logic and, through them, with 'language' in general. Vernacular languages such as English or French were considered messy and defective. It is only in the middle of the twentieth-century that things began to change, and that a descriptive attitude was adopted toward natural language within the literalist tradition.² Before that change occurred, context-sensitivity was taken to be a *defect* of natural language, like ambiguity. The fact that natural language sentences are indexical and therefore carry content only when uttered could therefore be deliberately ignored. Let us refer to this view (or rather, this attitude) as 'Proto-Literalism'.

Next in the development of the literalist tradition came 'Eternalism'. In contrast to Proto-Literalism, Eternalism was a substantial view regarding the phenomenon of indexicality in natural language. Indexicality was regarded as *not essential from a theoretical standpoint*. It was so considered because the following principle was widely accepted:

Eternalization Principle

For every statement that can be made in a natural language using a context-sensitive sentence in a given context, there is an eternal sentence, in that language (or in a suitable extension of that language), which can be used to make the same statement in any context.³

Thus indexicality turns out to be eliminable. Were it not for the necessities of practical life, we might utter only eternal sentences.

The Eternalization Principle has progressively been abandoned. It is now more or less accepted that natural language sentences are *irreducibly* context-sensitive. Some theorists even doubt the existence of eternal sentences in natural language.

² See e.g. Reichenbach (1947), Bar-Hillel (1954).

³ To obtain an eternal sentence from a context-sensitive one, one has only to replace the indexical constituents of the latter by non-indexical constituents with the same semantic value.

Eternalism, therefore, is out. But there still are fallback positions for Literalism. Indeed Literalism has been maintained, in progressively weaker forms, until today.

The strongest fallback position for Literalism consists in acknowledging the extent (and ineliminability) of context-dependence, while insisting that it still is the *sentence* which, in virtue of the rules of the language, expresses a content in context. This semantic notion of the content of a sentence (with respect to context) is held to be distinct from the pragmatic notion of the content of a speech act. For it is the linguistic conventions, not the speaker's intentions (or the hearer's beliefs regarding the speaker's intentions), which fix the content of the sentence with respect to context. Hence the name 'Conventionalism' for the view that *the truth-conditions of a sentence are fixed by the rules of the language independently of pragmatic considerations*. What determines the content of an indexical expression is not what is in the head of the language users, but a linguistic rule — the rule which constitutes the conventional meaning of that expression. As Barwise and Perry write, 'even if I am fully convinced that I am Napoleon, my use of « I » designates me, not him. Similarly, I may be fully convinced that it is 1789, but it does not make my use of « now » about a time in 1789' (Barwise and Perry, 1983: 148). It can therefore be maintained that natural language sentences possess a content (with respect to context) independently of the speech act which it is used to perform. The content of the speech act arguably depends upon the communicative intentions of the speaker which the utterance makes manifest to the hearer; but the content of the sentence is fixed directly by the rules of the language — with respect to context, admittedly, but independently of both the speaker's intentions and their recognition by the hearer.

Conventionalism replaced Eternalism when the Eternalization Principle was abandoned; and it still has advocates today. But Conventionalism is no longer the dominant position. It is widely acknowledged that the speaker's meaning has a role to play in fixing the truth-conditions of indexical sentences. To be sure, the reference of a pure indexical like 'I' is determined by a linguistic rule: the rule that 'I' refers to the speaker. But the reference of a demonstrative is not determined by a rule in this manner. It is generally assumed that there is such a rule, namely the rule that the demonstrative refers to the object which happens to be demonstrated or which happens to be the most salient, in the context at hand. But the notions of 'demonstration' and 'salience' are pragmatic notions in disguise. Ultimately, a demonstrative refers to *what the speaker who uses it refers to by using it*. Semantic

reference turns out to be parasitic on speaker's reference here. Even expressions like 'here' and 'now' which Kaplan classifies as *pure* indexicals (as opposed to demonstratives) are highly sensitive to the speaker's intent. The alleged rule of reference which is said to govern them is the rule that they refer to be the time or place of the context respectively; but what counts as the time and place of the context? How inclusive must the time or place in question be? It depends on what the speaker means, so that determining the content of words like 'here' and 'now' ultimately is a matter of pragmatics.

The alleged automaticity of content-determination and its independence from pragmatic considerations is an illusion due to an excessive concern with a sub-class of 'pure indexicals', namely words such as 'I', 'today', etc. In most cases, however, the reference of a context-sensitive expression is determined on a pragmatic basis.⁴ That is true not only of standard indexical expressions, but also of many constructions involving something like a free variable. For example, a possessive phrase such as 'John's car' arguably means something like *the car that bears relation R to John*. The free variable 'R' must be contextually assigned a particular value; but that value is not determined by a rule. What a given occurrence of the phrase 'John's car' means ultimately depends upon what the speaker who utters it means. That dependence upon speaker's meaning is a characteristic feature of semantically underdeterminate expressions, which are pervasive in natural language. Their semantic value varies from occurrence to occurrence, yet it varies not as a function of some objective feature of the situation of utterance but as a function of *what the speaker means*.

So we cannot maintain that the content of the sentence is fixed in context by linguistic rules. We must acknowledge the role of pragmatic considerations in determining truth-conditional content. This means that we must depart from Conventionalism; but there still is an ultimate fallback position for Literalism. According to that position, which I call 'Minimalism', the appeal to speaker's meaning in determining truth-conditional content is not free and unconstrained, but regulated by linguistic conventions. We appeal to speaker's meaning only when there is, in the meaning of the sentence type, a 'slot' to be filled pragmatically.

In the minimalist framework, the semantic content of the utterance departs

⁴ Thus John Perry (1997 : 595-6) distinguishes between 'automatic' indexicals and 'intentional' indexicals.

only minimally from the linguistic meaning of the sentence type (hence the name 'Minimalism'); it departs from it only when the meaning of the sentence itself requires that some contextual value be assigned to a context-sensitive word or morpheme, or to a free variable in logical form. The contextual assignment of values to indexicals and free variables is allowed to affect semantic content, because it is a bottom-up, linguistically controlled pragmatic process, i.e. a pragmatic process triggered (and made obligatory) by a linguistic expression in the sentence itself. But no other contextual influence is allowed to affect semantic content. In particular, 'top-down' pragmatic processes are banned. Such processes are not triggered by a particular expression in virtue of a linguistic rule, but take place in order to make sense of the speaker's communicative act in context. For example, sometimes we interpret what the speaker says nonliterally, because a literal interpretation would clash with the presumption that the speaker respects Grice's Cooperative Principle. Because they are not linguistically controlled, such interpretive processes have no impact on truth-conditions, according to Minimalism. They can only affect the overall content of the speech act performed by the speaker.

To sum up, four stages can be discerned in the historical development of the literalist tradition that started with Frege and is still dominant today. First came Proto-Literalism, according to which context-sensitivity is a defect of a natural language, to be ignored in theorizing about language. Then came Eternalism, which holds that indexicality is a practical convenience rather than an essential feature of natural language. Next came Conventionalism — the view that the conventional meaning of the sentence-type fully determines the content of the sentence (in context) independently of the speaker's meaning. Finally, Minimalism acknowledges the role of speaker's meaning in determining truth-conditions, but insists that the appeal to speaker's meaning is always subordinated to (controlled by) the conventional meaning of the sentence.

II. Towards Contextualism

As we have just seen, the strong forms of Literalism have been replaced by progressively weaker forms, in the historical development of the tradition stemming from Frege's work. How far will that process go? The currently accepted position is Minimalism. Will Minimalism be superseded by still weaker positions, and if so, at

what point will Literalism have to be squarely given up in favour of Contextualism? These are the questions I address in this section.

According to Minimalism, no contextual influences are allowed to affect the truth-conditional content of an utterance unless the sentence itself demands it. Yet, sometimes, the truth-conditions of an utterance seem to be affected by context in a top-down manner. For example, if I say 'It is raining', I mean that it is raining where I am (or at another contextually salient place), but nothing in the sentence seems to correspond to the place, which is provided by context without being linguistically 'articulated' (Perry, 1986). Faced with such cases, a defender of Minimalism has two options. He (or she) may bravely re-analyse the example so as to show that the pragmatic process at issue — here, the provision of a specific place — is a bottom-up process triggered by some expression in the sentence, appearances notwithstanding. Thus he may posit a free location variable in the logical form of the sentence (Stanley, 2000).⁵ Alternatively, the minimalist may draw a distinction between the semantic content of the sentence (here, the location-less proposition that it's raining at some place or other) and the content actually conveyed (viz. the proposition that it's raining where the speaker is). In contrast to the former, the latter need not obey the minimalist constraint.

The second of the two positions I have just described as available to the minimalist *concedes* that there are pragmatic processes that affect the interpretation of an utterance in a top-down manner, and that affect it at the level of (intuitive) truth-conditions. Therefore, by choosing this option rather than the first one, we move one step further in the direction of Contextualism. But we remain within the confines of Literalism because we maintain that the content of the sentence is the 'minimal' proposition determined by the linguistic meaning of the sentence when indexicals, free variables and other context-sensitive elements have been assigned contextual values. This position — which I call 'the Syncretic View' (Recanati, 2001, 2004) — is a compromise. On the one hand, the semantic content of the sentence is said to obey the minimalist constraint; on the other hand the intuitive content of the utterance can be freely enriched, as in this typical example from Scott Soames:

A man goes into a coffee shop and sits at the counter. The waitress asks him

⁵ This strategy defines the version of Minimalism which, in Recanati (2004), I call 'Indexicalism'.

what he wants. He says, "I would like coffee, please." The sentence uttered is unspecific in several respects — its semantic content does not indicate whether the coffee is to be in form of beans, grounds, or liquid, nor does it indicate whether the amount in question is a drop, a cup, a gallon, a sack, or a barrel. Nevertheless, it is obvious from the situation what the man has in mind, and the waitress is in no doubt about what to do. She brings him a cup of freshly brewed coffee. If asked to describe the transaction, she might well say, "He ordered a cup of coffee" or "He said he wanted a cup of coffee", meaning, of course, the brewed, drinkable kind. In so doing, she would, quite correctly, be reporting the content of the man's order, or assertion, as going beyond the semantic content of the sentence he uttered. (Soames, 2002 : 78)

Free enrichment — the process responsible for making the interpretation of an utterance more specific than its literal interpretation (as when 'coffee' is contextually understood as *coffee of the brewed, drinkable kind*) — is a top-down, pragmatically controlled pragmatic process. Another process of the same sort, 'predicate transfer' (Nunberg, 1995), takes us from a certain property, conventionally expressed by some predicative expression, to a distinct property bearing a systematic relation to it. For example, in 'I am parked out back', the property that is literally encoded by the verb phrase is a property of cars (the property of being parked out back), but the property which the expression actually contributes to the (intuitive) truth-conditions in this utterance is not a property of cars but another, systematically related property, namely the property a car-owner has when his or her car has the former property. In an utterance such as 'I am parked out back', transfer takes place because there is a linguistic mismatch between the predicate (which denotes a property of cars) and what it is applied to (a person). But such mismatch is not necessary for predicate transfer. Just as, through transfer, 'The ham sandwich left without paying' is understood as saying something about the customer who ordered the sandwich, 'The ham sandwich stinks' can be so understood, in a suitable context, even though the property of stinking potentially applies to sandwiches as well as to customers. Like free enrichment, the process of transfer is not a linguistically controlled but a pragmatically controlled pragmatic process: it is not triggered by something linguistic — some aspect of the linguistic signal being processed — but takes place in order to make sense of the communicative act performed by the speaker.

Predicate transfer and free enrichment are only two among a family of top-down pragmatic processes that affect the intuitive truth-conditions of utterances. This family of processes I call 'modulation', as opposed to the (bottom-up) process of assigning contextual values to indexicals, free variables etc. (Recanati, 2004). The Syncretic View acknowledges modulation, but limits its effects to the intuitive content of the utterance, that is, to the content of the speech act performed by the speaker. The content of the *sentence* (the 'minimal proposition' it expresses) is said to be unaffected, in accordance with Minimalism. But the Syncretic View can be criticized, on the grounds that the 'minimal proposition' it posits has no useful work to do. It is supposed to give us the semantic content of the sentence (as opposed to the content of the speech act), but do we really need to posit such a level of semantic content for the global sentence? Maybe we don't. What must ultimately be accounted for is what speakers say in the pragmatic sense — the content of their assertions (or of whatever speech acts they perform by their utterances). The job of linguistic meanings, semantic contents etc. is to contribute to the overall explanation. But, one may argue, it is sufficient to assign semantic contents (in context) to simple expressions. Modulation will operate on those contents, and the composition rules will compose the resulting senses, thereby yielding the content of the speaker's assertion. Of course it is possible to let the composition rules compose the plain semantic contents of the constituent expressions, thereby yielding the minimal proposition expressed by the sentence (an absurd proposition, in many cases). However, the content of the speaker's assertion will still be determined by composing the modulated senses resulting from the operation of pragmatic processes on the contents of the constituent expressions; so it is unclear what additional job the minimal proposition is supposed to be doing.⁶

Many people think that we need the minimal proposition because it is the input to the pragmatic processes which take us from what the speaker literally says to what she actually conveys. Those processes are said to operate globally on the output of the grammar. But that view has been rightly criticized. In 'There is a lion in the courtyard', 'lion' can be understood, through transfer, in the representational sense:

⁶ King and Stanley (forthcoming) offer an analogous argument purporting to show that it is fruitless to ascribe functional 'characters' to sentences: it is sufficient to ascribe characters to the parts, and redundant to ascribe characters also to the whole.

the thing that is said to be in the courtyard is not a (real) lion but a *representation* (more specifically, a statue) *of* lion. Now consider 'There is a stone lion in the courtyard'. What is said to be made of stone here? Clearly, it is the statue, rather than the lion which the statue represents. This simple fact shows that the process of representational transfer must take place *before* the composition rule associated with the noun-noun construction applies to the semantic values of the nouns 'stone' and 'lion'.⁷ If predicate transfer applied globally, after the grammatically triggered composition rules have applied, the interpretation we would get for the noun-phrase 'a stone lion' would be something like: *a representation of (a lion that is made of stone)*. But the correct interpretation is: *(a representation of a lion) that is made of stone*. We must therefore give up the Gricean idea that pragmatic processes operate globally on the output of the grammar.⁸ And this means that we don't really need the 'minimal proposition'.

The position I have just described I call 'Quasi-Contextualism'. It is very close to full-fledged Contextualism, but to get to the latter we need to take one more step.

So far we have granted that the pragmatic processes involved in modulation (free enrichment, transfer, etc.) are optional. For example, nothing prevents the sentence 'There is a lion in the courtyard' from being understood literally, as talking about a real lion. Or consider the following instance of free enrichment:

She took out her key and opened the door

The pragmatic process that enriches the meaning of this sentence so as to convey both a sense of temporal order (giving to 'and' the sense of 'and then') and a notion of the instrument used in opening the door (giving to 'opened the door' the sense of 'opened the door with the key') — that process might also not take place. As Grice

⁷ Note that this composition rule itself is context-sensitive (Partee, 1984: 294-5). The denotation of the compound results from intersecting the (literal, or pragmatically derived) denotation of the head noun with the set of objects that bear a certain relation *R* to the (literal, or pragmatically derived) denotation of the modifying noun. That relation can only be contextually determined. In 'stone lion', '*R*' is typically assigned the relation *being made of*, but in less accessible contexts a different relation will be assigned to the variable.

⁸ See Sag (1981), Recanati (1995), and Jackendoff (1997: 55 and 65-6).

emphasized, such pragmatic suggestions are always cancellable, explicitly or contextually. Once the pragmatic suggestion has been cancelled, what the words contribute to truth-conditional content is their bare linguistic senses.

From the optional character of modulation, it follows that the minimal proposition, even if it plays no causal-explanatory role, has at least the following, counterfactual status: it is the proposition which the utterance *would* express if no pragmatic process of modulation took place (Recanati, 1993: 318). To get full-fledged Contextualism we must deprive the minimal proposition even of this counterfactual status. While Quasi-Contextualism considers the minimal proposition as a theoretically useless entity, and denies that it plays any effective role in communication, Contextualism goes much further: it denies that the notion even makes sense. Contextualism ascribes to modulation a form of necessity which makes it ineliminable. *Without contextual modulation, no proposition could be expressed* — that is the gist of Contextualism. In this framework the notion of a 'minimal' proposition collapses: there *is* no proposition that is expressed in a purely 'bottom-up' manner.

III. Pragmatic Composition

To say that the pragmatic processes of modulation are optional is to say that in a suitable context, the senses expressed by the words would be, simply, the senses they possess in virtue of the rules of the language. The first of the three contextualist positions I am about to discuss — the Pragmatic Composition view (PC) — accepts that the literal, input sense undergoing modulation could, in a suitable context, be the expressed sense. So it construes the pragmatic processes of modulation as optional. But it construes them as optional only *with respect to the word whose sense is modulated*. If we consider not words in isolation, but the complex expressions in which they occur, we see that the pragmatic processes of modulation are not always contingent and dispensable, but often essential. Even though the linguistic meaning of a given word (or the semantic content we get after indexical resolution) could be the expressed sense, still the process of semantic composition, i.e. the putting together of that sense with the senses of other expressions, cannot proceed unless appropriate adjustments take place so as to make the parts fit together within an appropriate whole. On this view words have meanings which could go directly into

the interpretation, without modulation, but it is the composition process that forces modulation to take place, or at least invites it: often the meanings of individual words do not cohere by themselves, and can be fitted together only by undergoing a process of mutual adjustment.

Let us start with a simple example in which modulation is required to overcome a semantic mismatch.⁹

John hears the piano.

The verb 'hear' arguably denotes a relation between sentient organisms and sounds. Only sounds can be heard. Since a piano is not a sound, but a musical instrument, some adjustment is needed to make sense of 'hear the piano': either the noun-phrase 'the piano' must be given a metonymical interpretation, so that it stands for the sounds emitted by the piano; or (more plausibly) the verb 'hear' itself must be understood, not in its basic sense, but in a derived sense resulting from semantic transfer. An object is heard in the derived sense whenever the sound it emits is heard in the literal, basic sense.

We need to adjust or modulate the meaning of words even in the absence of linguistic mismatch. Think of an example like

John hates the piano.

A piano is certainly an object that *can* be hated, however strictly one construes the predicate 'hate'. Still, some contextual enrichment is in order, because to hate the piano is to hate it *under some aspect or dimension*. One may hate the sounds emitted by the piano, or one can hate playing the piano, or one can hate the piano as a piece of furniture, etc. The relevant dimension is contextually provided through the process of enrichment. (Similarly, if I say that Jim likes John's sister, the sense of 'like' will be — defeasibly — modulated so as to mean something different from what it means in 'Jim likes pork'.)

The crucial question is whether the sentence expresses a proposition independently of this type of modulation. To address this issue, let us consider

⁹ This example is borrowed from Langacker (1991 : 193-6).

another example, due to John Searle. The word 'cut' is not ambiguous, Searle says, yet it makes quite different contributions to the truth-conditions of the utterance in 'Bill cut the grass' and 'Sally cut the cake'. That is because background assumptions play a role in fixing satisfaction-conditions for the verb-phrase, and different background assumptions underlie the use of 'cut' in connection with grass and cakes respectively. We (defeasibly) assume that grass is cut in a certain way, and cakes in another way. Through enrichment the assumed way of cutting finds its way into the utterance's truth-conditions:

Though the occurrence of the word "cut" is literal in [both] utterances..., and though the word is not ambiguous, it determines different sets of truth conditions for the different sentences. 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. That is not what the speaker meant by his literal and serious utterance of the sentence.
(Searle, 1980 : 222-3)

Now an advocate of the Syncretic View will insist that a sentence such as 'Cut the grass' expresses something that has literal conditions of satisfaction quite independent of any background assumption; something very abstract, involving the constant meaning of 'cut' and not the variable senses it takes on particular uses (or types of use). Stabbing the grass with a knife and running over it with a lawnmower are two ways of literally obeying the order 'Cut the grass', on this view. But the contextualist remains skeptical. To get something genuinely evaluable, he claims, i.e. something which enables us to partition possible worlds into those in which the relevant condition is satisfied and those in which it is not, we need background assumptions (Searle, 1978). We cannot specify a determinate proposition which the sentence can be said literally to express, without building unarticulated assumptions into that proposition. The best we can do is to construct a disjunction of the propositions which could be determinately expressed by that sentence against alternative background assumptions.

In support of this controversial claim, Searle (1980) sets up an example for which no background assumption is readily available: 'Cut the sun'. What counts as obeying that order? We don't quite know. The abstract condition we can associate with that sentence (involving some form of linear separation affecting the integrity of the sun) is, precisely, too abstract to enable us to tell the worlds in which the condition is satisfied from the worlds in which it is not. It is not determinate enough to give us specific truth-conditions or obedience-conditions.

In previous writings I gave a real-life example of the phenomenon Searle is drawing our attention to.¹⁰ Consider the following dialogue from *Desire*, a film by Frank Borsage (1936):

- Pedro!
- Yes sir.
- Take the plate to the kitchen and disarm the fricassee.

What does the complex phrase 'disarm the fricassee' literally mean? It is hard to tell, even though we know the meanings of all the constituents. To make sense of that phrase, we must know the context. In the film, the context is as follows: (i) Gary Cooper (the speaker) is handing a fricassee plate to the waiter (Pedro); (ii) the fricassee plate contains a gun; (iii) that gun has just fallen from the hands of someone during a brief fight around the dinner table. With respect to that situation, the phrase 'disarm the fricassee' makes sense: it means that the waiter is to remove the gun from the plate. Without a proper background, however, we no more know the obedience conditions of Cooper's utterance 'Disarm the fricassee' than we know the obedience conditions of 'Cut the sun'.

In these examples, composing the senses of the parts so as to get a coherent sense for the whole involves imagining (or retrieving from memory) a possible scenario in which the senses of the parts fit together. That imaginative exercise involves elaborating what the meanings of the words give us — going beyond that linguistic meaning and, for example, interpreting 'disarm' in the specific sense of 'take the gun out of' or 'remove the gun from'.

As we shall see, a more radical version of Contextualism denies that words

¹⁰ See Recanati (1997 : 120 ; 1999 : 162-3).

like 'cut' possess a determinate sense: the constant meaning of 'cut' is more like an abstract schema which has to be fleshed out in context, and that is why elaboration is needed to get a determinate proposition. But the Pragmatic Composition view traces the need for modulation to the composition process, and some examples clearly support that view. Thus consider the adjective 'red'. Vagueness notwithstanding, it expresses a definite property: the property of being red or having the colour red. That property could, in principle, go into the interpretation of a sentence in which the adjective 'red' occurs. (For example: 'Imagine a red surface.')

But in most cases the following question will arise: what is it for the thing talked about to count as having that colour? *Unless that question is answered, the utterance ascribing redness to the thing talked about (John's car, say) will not be truth-evaluable.* It is not enough to know the colour that is in question (red) and the thing to which that colour is ascribed (John's car). To fix the utterance's truth-conditions, we need to know something more — something which the meanings of the words do not and cannot give us: we need to know *what it is for that thing (or for that sort of thing) to count as being that colour.* What is it for a car, a bird, a house, a pen, or a pair of shoes to count as red? To answer such questions, we need to appeal to background assumptions and world knowledge.¹¹ Linguistic competence does not suffice: pragmatic fine-tuning is called

¹¹ 'For a bird to be red (in the normal case), it should have most of the surface of its body red, though not its beak, legs, eyes, and of course its inner organs. Furthermore, the red color should be the bird's natural color, since we normally regard a bird as being « really » red even if it is painted white all over. A kitchen table, on the other hand, is red even if it is only painted red, and even if its « natural » color underneath the paint is, say, white. Moreover, for a table to be red only its upper surface needs to be red, but not necessarily its legs and its bottom surface. Similarly, a red apple, as Quine pointed out, needs to be red only on the outside, but a red hat needs to be red only in its external upper surface, a red crystal is red both inside and outside, and a red watermelon is red only inside. For a book to be red is for its cover but not necessarily for its inner pages to be mostly red, while for a newspaper to be red is for all of its pages to be red. For a house to be red is for its outside walls, but not necessarily its roof (and windows and door) to be mostly red, while a red car must be red in its external surface including its roof (but not its windows, wheels, bumper, etc.). A red star only needs to appear red from the earth, a red glaze needs to be red only after it is fired, and a red mist or a red powder are red not simply inside or outside. A red pen need not even have any red part (the ink may turn red only when in contact

for.

To sum up, on the view (PC) under discussion, even if the semantic content of a word is fixed by language (and context, if the expression is indexical), composing it with the contents of other words often requires help from above. It is semantic composition which has a fundamentally pragmatic character. So there is a sense in which modulation is necessary, but that is not quite the sense in which indexical resolution is. With indexical resolution there is a semantic gap and an instruction to fill the gap — both the gap and the instruction being part of the linguistic meaning of the expression. With modulation, there need be no gap and there is no instruction to search for some contextual filler. The expression means something, and that meaning could go into the interpretation — so modulation is optional — but to determine a suitable sense for complex expressions, we need to go beyond the meaning of individual words and creatively enrich or otherwise adjust what we are given in virtue purely of linguistic meaning. We must go beyond linguistic meaning, without being linguistically instructed to do so, if we are to make sense of the utterance.

IV. Literalist Responses to the Contextualist Challenge

According to Emma Borg (and other defenders of the Syncretic View), the fact that we are unable to specify intuitive conditions of application for the predicate 'cut the sun' does not support the contextualist conclusion that sentences per se do not have truth-conditions. There is, she claims, a crucial difference between 'knowledge of truth-conditions and the knowledge that truth-conditions are satisfied' (Borg, forthcoming). We may know the obedience-conditions of 'Cut the sun' in a purely 'disquotational' manner (i.e. we may know that 'Cut the sun' is obeyed iff the addressee cuts the sun), without knowing *what counts as* cutting the sun, in the context at hand. So there is no reason to deny sentences genuine truth-conditions. The sentence 'Oscar cuts the sun' does possess truth-conditions; such truth-conditions are determined by a recursive truth-theory for the language, which issues

with the paper). In short, what counts for one type of thing to be red is not what counts for another.' (Lahav, 1989 : 264)

theorems such as '*Oscar cuts the sun* is true iff Oscar cuts the sun'. We know those truth-conditions provided we know the language. What we don't know, simply in virtue of knowing the language, is 'a method of verification for those truth-conditions' (*ibid.*). This, then, is the syncretist's ultimate reply to the contextualist. According to the syncretist, the contextualist is guilty of endorsing a form of (so-called) 'verificationism'.¹²

This move strikes me as an unacceptable weakening of the notion of truth-condition. The central idea of truth-conditional semantics (as opposed to mere 'translational semantics') is the idea that, via truth, we connect words and the world.¹³ If we know the truth-conditions of a sentence, we know *which state of affairs must hold for the sentence to be true*. T-sentences display knowledge of truth-conditions in that sense only if the right-hand-side of the biconditional is *used*, that is, only if the necessary and sufficient condition which it states is transparent to the utterer of the T-sentence. If I say '*Oscar cuts the sun* is true iff Oscar cuts the sun', without knowing what it is to 'cut the sun', then the T-sentence I utter no more counts as displaying knowledge of truth-conditions than if I utter it without knowing who Oscar is (i.e. if I use the name 'Oscar' deferentially, in such a way that the right hand side is not really *used*, but involves some kind of mention).¹⁴

One may doubt the feasibility of referential or truth-conditional semantics and defend translational semantics as a viable alternative. I have heard (or read)

¹² *Ibid.* The first occurrence of this line of reply to contextualism can be found in Marcelo Dascal's discussion of Searle's 'Literal Meaning' (Dascal, 1981 : 173-4). The most recent occurrence I have seen is in Cappelen and Lepore (forthcoming).

¹³ See Lewis (1970 : 18-19), Evans and McDowell (1976 : vii-xi).

¹⁴ As Harman pointed out, if pure disquotational knowledge counts as knowledge of truth-conditions (in a suitably weak sense), then knowledge of truth-conditions (in that sense) does not count as knowledge of meaning. 'There is a sense in which we can know the truth conditions of an English sentence without knowing the first thing about the meaning of the English sentence. To borrow David Wiggins's (1972) example, we might know that the sentence « All mimsy were the borogroves » is true if and only if all mimsy were the borogroves. However, in knowing this we would not know the first thing about the meaning of the sentence, « All mimsy were the borogroves ».' (Harman, 1999 : 196)

arguments to that effect. My point however is that *if* we stick to the standard truth-conditional project (as Davidsonians like Cappelen and Lepore surely ought to do) then we should not accept the syncretist's claim that we somehow know the truth-conditions of 'Harry cut the sun'. (For we don't.)

The contextualist challenge is likely to elicit another unsatisfactory response, this time from the 'indexicalist'. To each dimension of contextual elaboration, the indexicalist may argue, there corresponds a slot in logical form, which must be filled for the utterance to say something definite. To illustrate that point, let us consider another contextualist example from Searle (1983 : 145-7).

When we ask someone to open the door, the content of the request goes beyond what is linguistically encoded. Not only is it necessary for the addressee to identify the relevant door (i.e. to complete or otherwise enrich the incomplete definite description 'the door'). She must also determine in what sense the door must be 'opened'. Besides doors and windows, eyes and wounds can be opened. Now if the addressee 'opened' the door by making an incision in it with a scalpel, as when opening a wound, she would not have satisfied the request. Still, in a special context, it could be that the request to open the door must be satisfied precisely by incising it by means of a scalpel. The manner of opening is thus defeasibly indicated by context, it is not determinable on the basis of just the linguistic meaning of the sentence (including the direct object of the verb). To be sure, we can make it explicit in the sentence itself by introducing supplementary details, but each addition of this sort cannot fail to introduce other underdeterminacies. If, for example, we add that the door must be opened 'with a key', we don't specify whether the key must be inserted into the lock or rather used like an axe to break the door open (Searle 1992 : 182). However explicit the sentence, there will always be some aspect of truth-conditional content that is contextually determined without being explicitly articulated.

At this point, the imagined indexicalist response consists in saying that, like all verbs, 'open' (or 'disarm' or 'cut') is associated with a complex frame¹⁵, involving a certain number of argument roles: a location playing the role of INSIDE; another location operating as OUTSIDE; a BOUNDARY separating the two; a MOVING OBJECT liable to pass from inside to outside (or the other way round); an OBSTACLE, that is, an

¹⁵ The notion of frame which I am using is that elaborated by Fillmore in a series of papers. See Fillmore (1976, 1982, 1985), and Fillmore and Atkins (1992).

entity preventing the passage of the moving object; an AGENT liable to free the passage by means of ACTION on the obstacle; an INSTRUMENT serving to accomplish the action; and so on and so forth. In context, each of the variables I have enumerated must be assigned a particular value: the INSIDE, the OUTSIDE, the OBSTACLE, the PATH etc., all must be contextually identified. In the case of 'opening a wound', the INSIDE is the interior of the wound, the OUTSIDE is the exterior of the body, the MOVING OBJECT are the internal secretions of the wound and so on. This contextual assignment of values to the variables is what determines the specific interpretation given to 'open' in a particular context, and it is no different from what is required for interpreting a context-sensitive expression. It is therefore unnecessary to modify semantic theory in order to give an account of Searle's examples; it is enough to extend the list of context-sensitive expressions, so as to include all verbs (insofar as they are all associated with frames which comprise a number of argument roles, the fillers of which must be contextually assigned).

This indexicalist response is no more convincing than the syncretist response was. Let's admit that the verb 'to open' is associated with the complex frame I have mentioned. Does that make it an indexical or context-sensitive expression, whose use triggers, indeed mandates, a contextual process of value assignment? No. There is an important difference between the argument roles of a frame and the indexical variables associated with context-sensitive expressions. Indexical variables *must* be contextually assigned values for the expression to acquire a definite semantic content. If the referent of 'he' in 'He boarded John's boat' is not contextually specified, or if the relation between John and the boat remains indefinite, the utterance does not have definite truth conditions. In contrast, the argument roles of a frame may but need not be assigned contextual values. The contextual assignment process is optional; it may, or may not, take place, depending on what is contextually relevant. In other words, it is the context (not the sentence) which determines which, among the many argument roles of a given frame, are contextually assigned particular values, and which remain indefinite (existentially quantified). In many contexts, it is of no importance whether the door is opened with a key or in another way; what counts is simply that it is opened. To be sure, for any given verb (or verb plus syntactic context), there is a small number of argument roles in the frame for which the contextual assignment of value is linguistically mandated; but the indexicalist response presupposes something much stronger: that the verb 'open' is like an

indexical expression, which acquires a definite content only when the argument roles of the associated frame (*all* the argument roles, insofar as they can all be contextually foregrounded) are contextually assigned values. That is evidently too strong. In a given context, many of the argument roles which feature in the frame are existentially quantified rather than contextually assigned values. This does not prevent the verb 'open' from expressing a definite content, in such a context.

To sum up, for indexicals it is the conventional meaning of the expression which triggers the process of indexical resolution and makes it mandatory. With ordinary expressions such as 'open', it is the context, not the conventional meaning of the expression, which is responsible for foregrounding certain aspects of the described situation and triggering a process of contextual specification which goes well beyond what is linguistically encoded. The process in question is top-down, not bottom up. It is a pragmatically controlled pragmatic process, rather than a linguistically controlled pragmatic process, like indexical resolution.

V. Radical Contextualism

PC is not the only possible contextualist position. According to another one — the Wrong Format view (WF) — it is not just semantic composition which requires adjustment and modulation of word meaning. Individual word meanings themselves are such that they *could not* go directly into the interpretation. They do not have the proper format for that. They are either too abstract and schematic, in such a way that elaboration or fleshing out is needed to reach a determinate content; or they are too rich and must undergo 'feature-cancellation', or some other screening process through which some aspects will be backgrounded and others focussed on. Note that there are versions of this view which take the meaning of a word to consist both in some abstract schema in need of elaboration *and* a large store of encyclopedic representations most of which must be screened off as irrelevant, on any particular use.

WF is more radical than PC, but a third contextualist position, Meaning Eliminativism (ME), is by far the most radical: it is a sort of WF pushed to the extremes. That position comes close to what I think Austin and Wittgenstein had in mind. Let me introduce it by contrasting it with WF.

According to WF, the sense expressed by an expression must always be

contextually constructed on the basis of the (overly rich or overly abstract) meaning, or semantic potential, of the word type. Just as the reference of an indexical expression is not linguistically given but must be contextually determined, the sense of an ordinary expression is not linguistically given but must be constructed. In that framework there still is a role for the linguistic meaning of word types: it is the input (or part of the input) to the construction process.

The difference between Meaning Eliminativism (ME) and WF is that, according to ME, we don't need linguistic meanings even to serve as input to the construction process. The senses that are the words' contributions to contents are constructed, but the construction can proceed without the help of conventional, context-independent word meanings.

Note that, according to a trivial extension of WF, the linguistic meaning of a word is not merely the input to the process of semantic modulation: it is also the output of a process of induction through which the child, or anyone learning the language, abstracts the meaning of the word from the specific senses which it expresses, or seems to express, on the observed occasions of use. It is a truism that the child or language learner starts not with pre-formatted linguistic meanings, but with actual uses of words and the contextualised senses that words assume on such uses. So both contextualised senses and context-independent linguistic meanings are input, and both are output, in some construction process. The linguistic meaning of a word type is the output of an abstraction process; that process takes as input the contextualised senses used as evidence by the language learner. On the other hand, the linguistic meaning of a word type also serves as input to the modulation process which yields as output the contextualised sense of the word on a particular occasion of use (Figure 1).


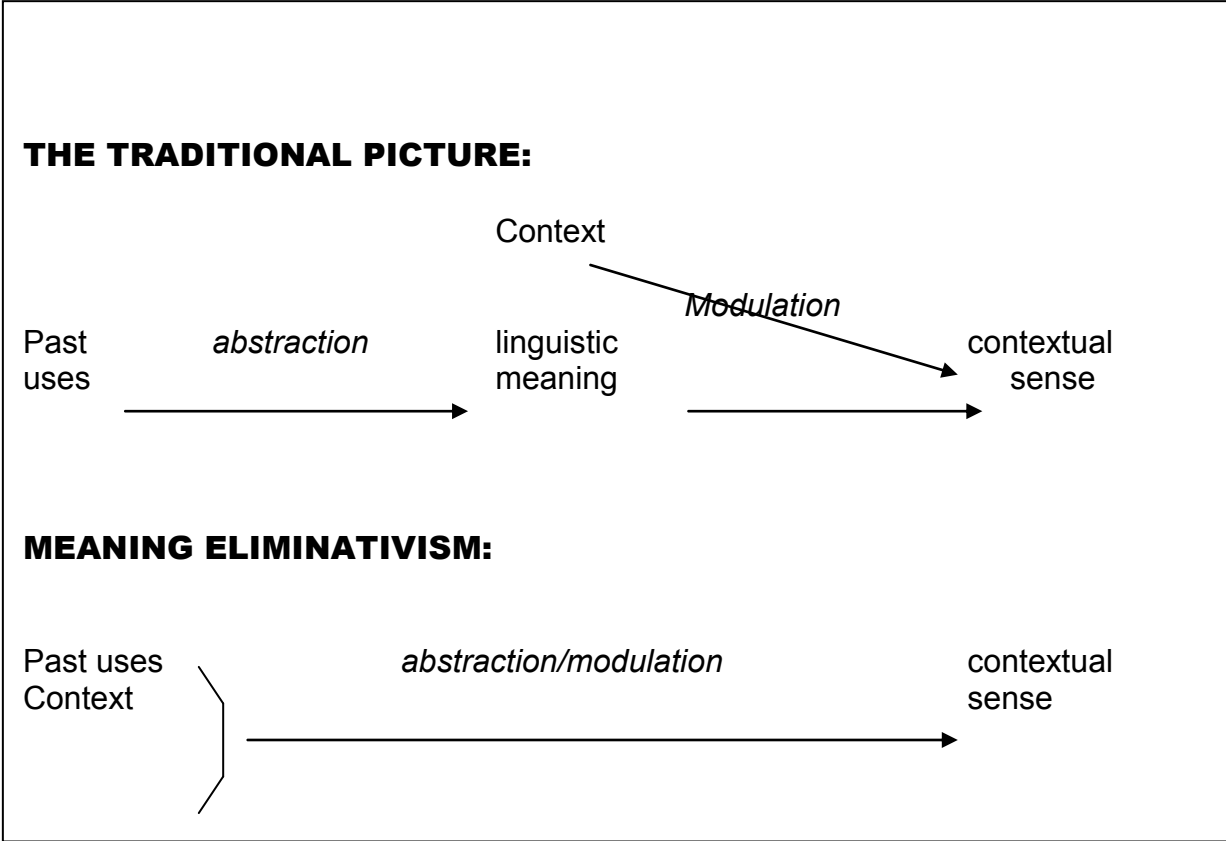


Figure 1

ME purports to simplify WF by *suppressing the intermediary step* (linguistic meaning) and computing directly the contextual sense which an expression assumes

on a particular occasion of use on the basis of the contextual senses which that expression had on previous occasions of use — without ever abstracting, or needing to abstract, 'the' linguistic meaning of the expression type.¹⁶ This amounts to merging the two construction processes: the abstraction of meaning from use, and the modulation of meaning in use (Figure 2). According to ME, there is a single process of abstraction-modulation which takes as input previous uses of the expression and yields as output the contextual sense assumed by the expression on the current use.



On the resulting picture, words are not primitively associated with abstract 'conditions of application', constituting their conventional meaning (as on the Fregean picture). The conditions of application for words must be contextually determined, like the reference of indexicals. What words, qua linguistic types, are associated with are not abstract conditions of application, but rather *particular applications*.

In the spirit of Wittgenstein, consider what it is for someone to learn a predicate P. The learner, whom I'll call Tom, observes the application of P in a particular situation S₁; he associates P and S₁. At this stage, the semantic potential of P for Tom is the fact that P is applicable to S₁. In a new situation S₂, Tom will judge

¹⁶ For a detailed psychological model supporting ME, see Hintzman (1986, 1988).

that P applies only if he finds that S_2 sufficiently resembles S_1 . To be sure, it is possible for S_2 to resemble S_1 in a way that is not pertinent for the application of P. The application of P to S_2 will then be judged faulty by the community, who will correct Tom. The learning phase for Tom consists in noting a sufficient number of situations which, like S_1 , legitimate the application of P, as opposed to those, like S_2 , which do not legitimate it. The semantic potential of P for Tom at the end of his learning phase can thus be thought of as *a collection of legitimate situations of application*; that is, a collection of situations such that the members of the community agree that P applies in or to those situations. The situations in question are the *source-situations*. The future applications of P will be underpinned, in Tom's usage, by the judgement that the situation of application (or *target-situation*) is similar to the source-situations.

In this theory the semantic potential of P is *a collection of source-situations*, and the conditions of application of P in a given use, involving a given target-situation S_3 , are *a set of features which S_3 must possess to be similar to the source-situations*. The set of features in question, and so the conditions of application for P, will not be the same for all uses; it is going to depend, among other things, on the target-situation. One target-situation can be similar to the source-situations in certain respects and another target-situation can be similar to them in different respects. But the contextual variability of the conditions of application does not end there. Even when the target-situation is fixed, the relevant dimensions for evaluating the similarity between that situation and the source-situations remain underdetermined: those dimensions will vary as a function of the subject of conversation, the concerns of the speech participants, etc.

One particularly important factor in the contextual variation is the relevant 'contrast set'. As Tversky (1977) has pointed out, judgements of similarity are very much affected by variations along that dimension. If we ask which country, Sweden or Hungary, most resembles Austria (without specifying the relevant dimension of similarity), the answer will depend on the set of countries considered. If that set includes not just Sweden, Hungary and Austria but also Poland, then Sweden will be judged more like Austria than Hungary; but if the last of the four countries considered is Norway and not Poland, then it is Hungary which will be judged more like Austria than Sweden. The explanation for that fact is simple. Poland and Hungary have certain salient geopolitical features in common which can serve as basis for the

classification: Hungary and Poland are then put together and opposed to Austria and Sweden. If we replace Poland by Norway in the contrast set a new principle of classification emerges, based on the salient features shared by Norway and Sweden: in this new classification Hungary and Austria go together. Tversky concludes that judgements of similarity appeal to features having a high 'diagnostic value' (or classificatory significance), and that the diagnostic value of features itself depends on the available contrast set.

So the set of similarity features on which sense depends itself depends upon the relevant contrast set, and the relevant contrast set depends upon the current interests of the conversational participants. It follows that one can, by simply shifting the background interests ascribed to the conversational participants, change the truth-conditions of a given utterance, even though the facts (including the target-situation) don't change, and the semantic values of indexicals remain fixed. Charles Travis has produced dozens of examples of this phenomenon of truth-conditional shiftiness over the last thirty years, and his examples often involve manipulating the relevant contrast-set.¹⁷

VII. Conclusion

In this paper I have discussed a number of positions, going from the early literalists'

¹⁷ See Travis (1975, 1981, 1989, 2000). The following example, inspired from Austin, is taken almost at random from a list of Travis-examples compiled by Claudia Bianchi (then a graduate student of mine):

Fred is walking with his young nephew beside a pond where a decoy duck is floating. Pointing to the decoy, he says, 'That's a duck'. Again we might ask whether what he said is true or false. But again, the above description is not enough for us to tell. If Fred has just finished laughing at a sportsman who blasted a decoy out of the pond, and if he has been trying to show his nephew how to avoid similar mistakes, then what he said is false. But suppose that Fred and his nephew are attending the annual national decoy exhibition, and the boy has been having trouble distinguishing ducks from geese. Then what Fred said may well be true. It would also be true had Fred said what he did in pointing out the fact that all the other ducks were poor copies (perhaps on the order of Donald Duck). (Travis, 1975 : 51)

blatant underestimation of context-sensitivity, to the most radical form of Contextualism. The positions are:

- Proto-Literalism
- Eternalism
- Conventionalism
- Indexicalism
- The Syncretic View
- Quasi-Contextualism
- PC
- WF
- Meaning Eliminativism

The first four positions stand squarely on the literalist side; the last three, squarely on the contextualist side. The Syncretic View and Quasi-Contextualism fall in between.

Literalism, in general, minimizes context-sensitivity. It strives at preserving the view that the proposition expressed by a (complete) sentence *is* the linguistic meaning of that sentence — or one of its meanings, if the sentence is ambiguous. The only exception that is allowed for is indexicality, which is not considered as a threat to the general picture because it is a form of context-sensitivity which remains under linguistic control. Indexicalism goes as far as to generalize indexicality in order to protect semantic content from ‘top-down’ or ‘strong’ pragmatic effects — a form of context-sensitivity that is *not* under linguistic control.

Such an exclusion of ‘top-down’ or ‘strong’ pragmatic effects on truth-conditions I find dogmatic. If we give up the stronger forms of Literalism and admit that the content of an utterance is not entirely fixed by linguistic rules, but has to be contextually determined by making sense of the speaker’s speech act, is it not obvious that some aspects of content may happen to be contributed entirely by context ? Why insist that *all* aspects of content must be traceable to aspects of linguistic form, if not because one is still in the grip of the literalist prejudice ?

Minimalism can be defended, by explicitly going stipulative. One may grant the existence, or at least the possibility, of strong pragmatic effects, while defining ‘the proposition literally expressed by an utterance’ in such a way that it can only satisfy the minimalist constraint. In other words, one may draw a distinction between what is

said in the intuitive sense — the actual content of one's utterance — and the proposition which can be assigned to that utterance as its 'literal' content, i.e. the minimal content that results from contextually assigning values to all indexical or free variables. That is the gist of the 'Syncretic View'. In this framework the proposition literally expressed satisfies Minimalism by definition : it does not incorporate the output of pragmatic processes unless they are mandatory and triggered by elements in the syntactic structure of the sentence.

What is the point of positing such a minimal proposition? As I have emphasized, it is unclear that it plays any role in the actual process of interpretation. This much must be conceded to the quasi-contextualist. It has been argued that we need the minimal proposition to account for 'the character of the information available to the hearer' (Bach, 1994 : 158). The minimal proposition, Bach says, is 'included in the information available to the hearer in understanding an utterance' (Bach, 1994 : 159). What this means, presumably, is that the hearer knows the literal semantic values of the constituents, and knows the appropriate composition rules. He should therefore be credited with the ability to compose those values so as to determine the literal semantic value of the whole — the minimal proposition. In practice, that need not be done. Since modulation takes place locally, the interpreter does not actually compose the literal semantic values of the constituents to determine the minimal proposition ; rather, he directly determines what is said (in the intuitive sense) by composing the pragmatic values resulting from whatever pragmatic processes locally operate on the literal semantic values of the constituents. Be that as it may, the minimal proposition is said to be 'available to the hearer, even if not actually accessed' (Bach, 1994 : 158). The interpreter does not compute it, but he could.

Full-fledged Contextualism questions the claim that, independent of modulation, it is possible to determine a minimal proposition by mechanically composing the meanings of the constituents. I have briefly indicated the sort of argument a contextualist may put forward in support of this denial, but the issue is far from being settled. My intention was not to argue for (or against) Contextualism in this paper, but only to convince you that the debate ought to take place. This means that we must get rid of the last literalist prejudice: we must stop *presupposing* that there is such a thing as the minimal proposition expressed by an utterance. It is important to realize that that literalist assumption, pervasive though it is among philosophers of language, rests on a substantial and highly controversial conception

of both word meaning and sentence meaning. There is no reason to rule out, a priori, a contextualist account of word and sentence meaning, even if such an account entails the nonexistence of 'minimal propositions'.

References

- Bach, Kent (1994). 'Conversational Implicature', in *Mind and Language*, 9 : 124-62.
- Bar-Hillel, Yehoshua (1954). 'Indexical Expressions', *Mind*, 63 : 359-79.
- Barwise, Jon and Perry, John (1983). *Situations and Attitudes* (Cambridge, Mass. : MIT Press).
- Borg, Emma (forthcoming). 'Saying What You Mean : Unarticulated Constituents and Communication'.
- Cappelen, Herman and Lepore, Ernie (forthcoming). 'Radical and Moderate Pragmatics : Does Meaning Determine Truth-Conditions ?'
- Cohen, Jonathan (1971). 'Some Remarks on Grice's Views About the Logical Particles of Natural Language', in Bar-Hillel, Yehoshua (ed.), *Pragmatics of Natural Languages* (Dordrecht : Reidel), 50-68.
- Cohen, Jonathan (1993). 'The Semantics of Metaphor', in Ortony, Andrew (ed.) *Metaphor and Thought* (2nd ed., Cambridge : Cambridge University Press), 58-70.
- Conant, J. (1998). 'Wittgenstein on Meaning and Use', *Philosophical Investigations*, 21: 232-250.
- Dascal, Marcelo (1981). 'Contextualism', in Parret, Herman, Sbisà, Marina and Verschueren, Jeff (eds.), *Possibilities and Limitations of Pragmatics* (Amsterdam : Benjamins), 153-77.
- Evans, Gareth and McDowell, John (1976). 'Introduction' to *Truth and Meaning : Essays in Semantics* (Oxford : Clarendon Press), vii-xxiii.
- Fillmore, Charles (1976). 'Frame Semantics and the Nature of Language', *Annals of the New York Academy of Sciences*, 280 : 20-32.
- Fillmore, Charles 'Frame Semantics', in The Linguistic Society of Korea (ed.) *Linguistics in the Morning Calm* (Seoul : Hanshin), 111-38.
- Fillmore, Charles (1982). 'Frames and the Semantics of Understanding', *Quaderni di Semantica*, 6 : 222-54.
- Fillmore, Charles and Atkins, Sue (1992). 'Toward a Frame-Based Lexicon: The Semantics of RISK and its Neighbors', in Lehrer, Adrienne and Kittay, Eva (eds.),

- Frames, Fields and Contrasts: New Essays in Semantic and Lexical Organization* (Hillsdale, N.J. : Lawrence Erlbaum Associates), 75-102.
- Grice, Paul (1989). *Studies in the Way of Words* (Cambridge, Mass. : Harvard University Press).
- Harman, Gilbert (1999). 'Meaning and Semantics', in his *Reasoning, Meaning, and Mind* (Oxford : Clarendon Press), 192-205.
- Hintzman, Douglas (1986). "'Schema Abstraction" in a Multiple-Trace Memory Model', *Psychological Review*, 93 : 411-28.
- Hintzman, Douglas (1988). 'Judgments of Frequency and Recognition Memory in a Multiple-Trace Memory Model', *Psychological Review*, 95 : 528-51.
- Jackendoff, Ray (1997). *The Architecture of the Language Faculty* (Cambridge, Mass. : MIT Press).
- Kaplan, David (1989). 'Demonstratives', in Almog, Joseph, Wettstein, Howard, and Perry, John (eds.), *Themes from Kaplan* (New York : Oxford University Press), 481-563.
- King, Jeffrey and Stanley, Jason (forthcoming). 'Semantics, Pragmatics, and the Role of Semantic Content'.
- Lahav, Ron (1989). 'Against Compositionality: the Case of Adjectives', *Philosophical Studies*, 57 : 261-79.
- Langacker, Ronald (1991). *Concept, Image, and Symbol : The Cognitive Basis of Grammar* (Berlin : Mouton/De Gruyter).
- Lewis, David (1970). 'General Semantics', *Synthese*, 22 : 18-67.
- Nunberg, Geoffrey (1995). 'Transfers of Meaning', *Journal of Semantics*, 12 : 109-32.
- Partee, Barbara (1984). 'Compositionality', in Landman, Fred, and Veltman, Frank (eds.), *Varieties of Formal Semantics* (Dordrecht : Foris), 281-312.
- Perry, John (1986). 'Thought Without Representation', *Proceedings of the Aristotelian Society, Supplementary Volume*, 60 : 137-52.
- Perry, John (1997). 'Indexicals and Demonstratives', in Hale, Bob and Wright, Crispin (eds.), *A Companion to the Philosophy of Language* (Oxford : Blackwell), 586-612.
- Recanati, François (1993). *Direct Reference: From Language to Thought* (Oxford : Blackwell).
- Recanati, François (1995). 'The Alleged Priority of Literal Interpretation', *Cognitive Science*, 19 : 207-32.
- Recanati, François (1997). 'La Polysémie Contre le Fixisme', *Langue Française*,

113 : 107-23.

Recanati, François (1999). 'Situations and the Structure of Content', in Murasugi, Kumiko, and Stainton, Rob (eds.), *Philosophy and Linguistics* (Boulder : Westview Press), 113-65.

Recanati, François (2001). 'What is Said', *Synthese*, 128 : 75-91.

Recanati, François (2004). *Literal Meaning* (Cambridge : Cambridge University Press).

Reichenbach, Hans (1947). *Elements of Symbolic Logic* (London : Macmillan).

Sag, Ivan (1981). 'Formal Semantics and Extralinguistic Contexts', in Cole, Peter (ed.), *Radical Pragmatics* (London : Academic Press), 273-94.

Searle, John (1978). 'Literal Meaning', *Erkenntnis*, 13 : 207-24.

Searle, John (1980). 'The Background of Meaning', in Searle, John, Kiefer, Ferenc, and Bierwisch, Manfred (eds.), *Speech Act Theory and Pragmatics* (Dordrecht : Reidel), 221-32.

Searle, John (1983). *Intentionality* (Cambridge : Cambridge University Press).

Searle, John (1992). *The Rediscovery of the Mind* (Cambridge, Mass. : MIT Press).

Soames, Scott (2002). *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity* (New York : Oxford University Press).

Stanley, Jason (2000). 'Context and Logical Form', *Linguistics and Philosophy*, 23 : 391-434.

Travis, Charles (1975). *Saying and Understanding* (Oxford : Blackwell).

Travis, Charles (1981). *The True and the False: : the Domain of the Pragmatic* (Amsterdam : Benjamins).

Travis, Charles (1989). *The Uses of Sense: Wittgenstein's Philosophy of Language* (Oxford : Clarendon Press).

Travis, Charles (2000). *Unshadowed Thought* (Cambridge, Mass. : Harvard University Press).

Tversky, Amos (1977). 'Features of Similarity', *Psychological Review*, 84 : 327-52.