

Pragmatics and Logical Form

François Recanati

Institut Jean-Nicod, Paris

1. Truth-conditional pragmatics

Robyn Carston and I, along with many others, share a general methodological position which I call 'Truth-Conditional Pragmatics' (TCP). TCP is the view that the effects of context on truth-conditional content need not be traceable to the linguistic material in the uttered sentence. Some effects of context on truth-conditional content 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' pragmatic processes that take place not because the linguistic material demands it, but because the utterance's content is not faithfully or wholly encoded in the uttered sentence, whose meaning requires adjustment or elaboration in order to determine an admissible content for the speaker's utterance.

The extra step required to get from conventional meaning to admissible content is usually treated as external to truth-conditional content proper, because truth-conditional content proper is supposed to be unaffected by pragmatic considerations *unless* such considerations are forced upon the interpreter by the linguistic material itself. Now we can perhaps characterize a notion of literal content such that literal content is, *by definition*, independent of pragmatic considerations (unless such considerations are imposed by the linguistic material itself), but when it comes to the *intuitive* truth-conditions of an utterance, TCP holds that they result, in part, from pragmatic processes that are not triggered by the linguistic material. Assuming that semantics is to account for the intuitive truth-conditions of utterances, it must make room for 'free' (pragmatically controlled) pragmatic processes, just as it makes room for linguistically controlled pragmatic processes in order to secure contextual values for the context-sensitive elements in the sentence.

Free pragmatic processes take as input the meaning which is the semantic interpretation of some expression and yield as output the modulated meaning that will undergo semantic composition with the meanings of the other expressions in the sentence. In other words, the composition rules determine the value of a complex expression on the basis of the *pragmatically modulated* values of the parts, according to formula (F):

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

In that formula ‘ I ’ stands for the interpretation function, ‘ a^b ’ stands for a complex expression formed from the parts ‘ a ’ and ‘ b ’, and the ‘ g ’s are free higher-order variables ranging over available pragmatic functions (including identity, which gives us the ‘literal’ case).¹ 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. Pragmatic modulation covers optional processes such as free enrichment, loosening, metonymic transfer, etc. : processes which (arguably) affect the intuitive truth-conditions but which take place for pragmatic reasons, without being triggered by the linguistic material in an obligatory manner.²

One way of understanding the formula is to say that *semantic composition itself is a context-dependent process* : in the course of deriving the semantic value of a complex expression, one optionally modulates the semantic values of the parts, and it is the context which determines which pragmatic function, if any, comes into play and yields the modulated value that undergoes semantic composition. This corresponds to the view which, in my book *Literal Meaning*, I called ‘Pragmatic Composition’.³ Another, even more radical way of understanding the formula corresponds to a view put forward by Gennaro Chierchia in

¹ I am indebted to Gennaro Chierchia for discussion of the overall picture which I am presenting here.

² As Peter Pagin pointed out to me (personal communication), this formula is insufficiently general. The format must be that of recursion over modulated meaning : modulation (possibly empty) occurs at each node, giving the (combined) modulated meaning at that node. Thus, with M for semantic meaning and Mod for modulated meaning, and g_i etc. for particular modulation functions, Pagin suggests that we have

$$\begin{aligned} Mod(e) &= g_i(M(e)) \text{ for some } g_i, \text{ in case } e \text{ is simple} \\ Mod(s(e_1, \dots, e_n)) &= g_i(c(Mod(e_1), \dots, Mod(e_n))), \text{ for some } g_i \end{aligned}$$

where s is the syntactic operation, and c the ordinary composition function.

³ Recanati 2004 : 138-40. See Jackendoff 1997 : 47-67 and 2002 : 387-94 for a similar notion of ‘enriched composition’, and Pagin (2005) for a general discussion of context-dependence and compositionality.

connection with scalar implicatures (Chierchia 2004). On Chierchia's picture, the interpretation 'function' is no longer a function but a relation. Adapting Chierchia's idea, we could say that each expression denotes a set of admissible values: the same linguistic form receives an indefinite number of distinct, alternative denotations, depending on which optional pragmatic processes (which 'g's distinct from identity) come into play. Thus 'tiger', in the right context, comes to mean 'representation of tiger', '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. That suggests that in those phrases, the words 'lion' and 'straight' get a modulated value, distinct from their standard semantic value.

Whichever construal we favour, it is important to realize that the variables over pragmatic functions that occur in formula (F) are there only in the theorist's metalanguage. They are not supposed to be present at any level of syntactic structure in the object-language. That is, indeed, what defines free pragmatic processes: they are not triggered by a variable in the syntax, or anything of the sort, but take place for purely pragmatic reasons — in order to adjust the conventional meaning of the words to the situation at hand. Even though they have an impact on truth-conditional content, they are a matter of *use*, not a matter of conventional meaning.

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, Szabo 2000, King and Stanley 2005) 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 there are theorists (e.g. Predelli 2005) who are suspicious of it because they take it to threaten the very enterprise of semantics. This seems to me grossly exaggerated. TCP may complicate the task of semantics but certainly does not make it impossible. Be that as it may, I will assume TCP in what follows, and will only be concerned with issues concerning its proper interpretation. What exactly does TCP say? How are we to understand free pragmatic processes? One particular answer, favoured by Robyn, is that which has been provided by Sperber and Wilson and constitutes a central tenet of relevance theory. In this paper, I will not directly argue for or against that answer; rather, I will place it among other possible answers so as to structure the theoretical space and make further discussion possible.

2. Free pragmatic processes: two interpretations

Although variables for pragmatic functions are confined to the metalanguage and are not syntactically projected, still there are two possible ways of looking at the role played by free pragmatic processes. One construal is ‘syntactic’ and the other one is ‘semantic’ (Recanati 2002: 339-42, Stanley 2005: 237). Relevance theorists opt for the syntactic construal, but it will be convenient to start with the other, semantic construal.

On the semantic construal, the output of free pragmatic processes is not a level of representation at all. It is a proposition, i.e. a semantic object. That proposition is the (intuitive) interpretation of the utterance, and it depends upon two things in addition to the semantic rules of the language: the *logical form* of the sentence serves as input to the interpretation process, while the *pragmatic context* determines both the semantic values of context-sensitive expressions in logical form and the pragmatic functions which optionally come into play in deriving the semantic value of the whole from the (possibly modulated) semantic values of the parts.

In contemporary generative linguistics, the logical form of a sentence, or LF, is standardly construed as a level of syntactic representation that is the proper input to semantic interpretation. At that level, important logical properties of the sentence such as the relative scope of quantifiers and anaphoric dependencies are formally displayed, in such a way that that level of syntactic representation can be systematically mapped to logical formulae which capture the inferential potential of the sentence (hence the name ‘logical form’ for that level of syntactic representation). The coexistence of syntactic structures of a certain sort and the logical representations associated with them under the heading ‘logical form’ creates a potential ambiguity, which Chierchia and McConnell-Ginet propose to avoid by distinguishing the logical form qua syntactic structure (LF) from the logical formula it maps to (*lf*).

Now the proposition which is the output of semantico-pragmatic processing itself determines a ‘logical form’ for the utterance i.e. a certain inferential profile: what the utterance entails and what it is entailed by, in virtue of the proposition it expresses. That ‘logical form’ too can be represented as a formula in a logical calculus — possibly the same calculus as that from which *lf*s are drawn. To bring out the analogy, I will use ‘*lf**’ as an abbreviation for the modified logical form that results from semantico-pragmatic processing, and which corresponds to the utterance’s intuitive truth-conditions (as opposed to its minimal

or literal truth-conditions, i.e. what we get if we submit LF to semantic interpretation without any recourse to free pragmatic processes).

On the semantic construal, the modified logical form *lf** is only a perspicuous representation, in a well-behaved logical language, of the utterance's intuitive truth-conditions, which truth-conditions result from interpreting the utterance's LF in accordance with formula (F) above. I call this construal 'semantic' rather than 'syntactic' because free pragmatic processes come into play purely as a matter of interpretation. They do not give rise to a further level of representation in addition to the syntactic representations (LFs) which serve as input to the interpretation process. In particular, *lf**s are not 'semantic representations' in the object-language. In the tradition of referential semantics (as opposed to translational semantics), interpretation proceeds by mapping representations (e.g. LFs) to worldly entities or complexes of such, not by mapping them to further representations.

I now turn to the 'syntactic' interpretation, favoured by relevance theorists and most researchers in pragmatics. The starting point is the same: there is a level of syntactic representation — LF — that incorporates « whatever features of sentences structure (1) enter directly into the semantic interpretation of sentences, and (2) are strictly determined by properties of sentence grammar » (Chomsky 1976 : 305). What distinguishes the syntactic from the semantic construal is the following claim, also made by Chomsky : Representations at LF are mapped to *more elaborate representations* « which may involve belief, expectations and so on in addition to properties of LF determined by grammatical rule » (*id.*). These additional representations are the modified logical forms (*lf**s) on the syntactic construal. Thus relevance theorists take pragmatic processing to operate on representations, *and to output further representations*. The representation operated on is the utterance's 'logical form', as delivered by the linguistic module. The representation which results from pragmatic processing is the modified logical form (*lf**), syntactically interpreted. Sperber and Wilson call it the utterance's 'propositional form'.

Is the syntactic construal consistent with referential semantics or does it necessarily go together with a translational approach to semantics ? This is an interesting and tricky issue. Robyn Carston suggests that relevance theory is inconsistent with truth-conditional semantics but thinks this is ok since « work in other frameworks shows... that giving an account of natural language semantics in terms of [conceptual representations] is very much a live option (see Katz 1972, Jackendoff 1983, 1990) » (Carston 2002 : 89-90). Like Sperber and Wilson, she opts for a translational approach, where semantics maps syntactic structures to conceptual structures (which can then be modified or 'developed' through pragmatic processing). I will

discuss the relevance-theoretic view in some detail below (§3), but for now I want to defend the following claim : the syntactic construal of TCP, by itself, is *not* inconsistent with the project of giving a truth-theoretic (= referential) semantics for natural language.

Following the majority view in semantics, let us assume that semantic interpretation is referential rather than translational : it maps representations to *what they represent*, rather than mapping them to further representations. This is in contrast to the ‘translational’ or ‘cognitive’ view, according to which semantics maps syntactic representations to ‘semantic representations’ construed as partial or schematic mental representations. Still, *nothing prevents a truth-conditional semanticist from bringing mental representations into the picture*. Indeed, the ‘logical forms’ which undergo semantic interpretation in the truth-theoretic framework can *themselves* be construed as mental representations. As Chierchia writes,

The hypothesis of a logical form onto which [surface] syntactic structure is mapped fits well with the idea that we are endowed with a language of thought, as our main medium for storing and retrieving information, reasoning, and so on. The reason why this is so is fairly apparent. Empirical features of languages lead linguists to detect the existence of a covert level of representation with the properties that the proponents of the language of thought hypothesis have argued for on the basis of independent considerations. It is highly tempting to speculate that logical form actually *is* the language of thought. (Chierchia 1999 : c-ci)

On this view, which has the merit of simplicity, logical forms are conceptual representations, yet conceptual representations that are strictly determined by the grammar (and as such belong to the language system). They belong both to the linguistic system and to the conceptual system, and serve as interface between the two systems. This view is consistent with mainstream referential semantics because such representations are construed as the *input* to semantic interpretation, in the truth-theoretic sense, rather than its output, as in translational/cognitive semantics (where semantics is said to map syntactic structures to conceptual representations).

In this framework pragmatic processing may be allowed to ‘elaborate’ the logical forms, qua conceptual representations, into further conceptual representations. Indeed, *lf**s can be construed as mental representations resulting from pragmatic operations on *lfs*. This is the syntactic construal of free pragmatic processes: they are seen as mapping the mental

representations associated with sentences in virtue of their grammatical properties onto further mental representations resulting in part from pragmatic processing.

3. Logical form in relevance theory

Relevance theorists think the project of giving a truth-conditional semantics for natural language is doomed to failure, and they assume a translational semantics à la Jerry Katz. Semantic interpretation is viewed as (part of) a ‘decoding’ process, through which syntactic representations are systematically mapped to semantic representations. Those semantic representations, resulting from decoding, are what Sperber and Wilson call ‘logical forms’. They are conceptual structures with logical properties (hence they can undergo logical operations), but they are distinct from thoughts in that they are not ‘complete’ and truth-evaluable. The logical forms that are associated with sentences as their meanings are partial (gappy) mental representations. They are comparable to (and can be represented by means of) predicate calculus formulae containing free variables (plus instructions on how to fill the variables).⁴

According to relevance theory, pragmatic processing takes us from the gappy mental representations associated with sentences in virtue of the semantic rules of the language to full-fledged mental representations, namely the thoughts which the contextualized utterances express. Those thoughts are semantically complete — truth-evaluable — and constitute the ‘propositional forms’ of the utterances which express them. Free pragmatic processes operate along the way : in the course of what Sperber and Wilson call the ‘development’ of the logical form, values are assigned to variables in the logical form, but enrichment or loosening take place as well. The whole process is syntactic in the following sense : it operates on representations and yields further representations. But only the input representation is ‘linguistic’. The output representation is a *mental* representation, corresponding to the contextual interpretation of the utterance. It is a sentence, but a sentence in the language of thought.

On the RT view, logical forms, the output of linguistic decoding, belong to the language system (they are the ‘semantic representations’ of translational semantics) but they are also conceptual : they are partial or schematic conceptual representations. So logical form is the interface between the language system and the thought system, on this view, as it is on

⁴ See Carston 2002 : 60 for an example.

Chierchia's. What, then, is the difference between Chierchia's view and the relevance-theoretic view ?

One superficial difference is that Sperber and Wilson take logical forms to be the output of semantic interpretation (in the translational sense), while Chierchia takes them to be the input to semantic interpretation (in the truth-theoretic sense). I call this a 'superficial' difference because there is no real conflict here, since two different notions of semantic interpretation are at stake. Nothing prevents a theorist from holding that logical forms are both the output of semantic interpretation in a first sense, and the input to semantic interpretation in a second sense. Robyn has sketched such an oecumenical view :

This position has been expressed often in the relevance-theoretic literature in talk of 'two types of semantics' : (1) A translational linguistic semantics, which could be described in statements of the form '*abc* means (= encodes) '*ijk*', where '*abc*' is a public-language form and '*ijk*' is a Mentalese form (most likely an incomplete, schematic Mentalese form) ; (2) a 'real' semantics, which explicates the relation between our mental representations and that which they represent (so it must be 'disquotational'). (Carston 2002 : 58)

But there is a deeper difference, and it concerns the proper input to 'real' semantics. On Chierchia's picture, as in mainstream generative linguistics more generally, the logical forms are a level of syntactic representation *which is semantically interpretable* (in the sense of real, truth-theoretic semantics). Logical forms are both syntactic representations delivered by the language system *and* conceptual representations that can be semantically evaluated (modulo saturation). They are representations endowed with « a logical syntax appropriate for recursively stating the truth-conditions of a sentence » (Hornstein 1995 : 5). Now this is something which relevance theorists do not accept. For relevance theorists, the linguistic representations which the linguistic module delivers are too indeterminate to be the input to semantic evaluation procedures. They must first be elaborated pragmatically. In other words, the two aspects which go together in the mainstream notion of logical form (being determined by grammar, being semantically evaluable) are disjoined in relevance theory: what the grammar delivers (the logical form in the sense of relevance theory) is not (yet) semantically evaluable. Apparently, 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).

At this point, one might think that even this ‘deeper’ difference is terminological. Neither referential values nor truth-values can be assigned to linguistic forms independent of context. That much is granted by everybody. So logical forms, qua properties of sentence-types, are not truth-evaluable. They are incomplete and ‘gappy’. If that is what relevance theorists mean when they say that logical forms cannot be given a ‘real’ semantics (and this is clearly *part* of what they mean), *then* there is no conflict with the truth-conditional view, contrary to what they assume. From the standpoint of truth-conditional semantics, logical forms are semantically interpretable only in the following sense : they can be assigned truth-conditions *relative to particular assignments of values to its context-sensitive elements*. But they cannot be assigned absolute truth-conditions (independent of context). So, there is a sense in which they are not semantically interpretable, but there is also a (weaker) sense in which they are.

Still, the disagreement between truth-conditional semantics and relevance theory is not merely terminological. Relevance theorists (like other TCP theorists such as Kent Bach) deny that logical forms are semantically interpretable *even in the weak sense*. What they call semantic under-determinacy goes beyond indexicality. For them, logical properties like quantifier scope and anaphoric dependencies are not fixed at the level of logical form (qua delivered by the grammar), but only at the further conceptual level of ‘propositional form’. Ruth Kempson concludes that

natural languages are not directly semantically (i.e. truth-theoretically) interpreted. Natural language expressions are provided an interpretation by processes of grammar only in the sense that they are associated with some construct in a system of representations which is said to constitute the language of thought. (...) This system of representations onto which natural language objects are mapped is itself a semantically transparent system, with a recursively definable truth-theoretic semantics for all expressions of the system, a principle of compositionality applying strictly to determine the semantic properties of all complex constituents on the basis of their parts. It is in this language-of-thought system that inference is definable, not in any natural-language grammar. (Kempson 1993 : 72-73)

Kempson sometimes makes this point by saying that LF (as traditionally conceived) is not a natural-language structure but a language-of-thought structure. The mapping from (surface) syntactic structures to LF is therefore not part of the grammar – though it is constrained by the grammar. All LF-building processes, she says, must be construed as filters on the pragmatic process of constructing a fully specified propositional representation (a pragmatic process that may well involve free pragmatic processes as sub-components).

4. *How many systems, and how different ?*

Since they are not strictly determined by the grammar, modified logical forms do not belong to the language system, but to a different ‘system of representation’, as Chomsky puts it : the conceptual system (Fodor’s ‘language of thought’). Now Jackendoff emphasizes the *heterogeneity* between the two systems of representation. According to Jackendoff, 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 disjoint, 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 (see Jackendoff 2002 : 270).

Other theorists have denied the alleged heterogeneity. For Chierchia, as we have seen, logical forms are *already* representations in the language of thought — they are conceptual representations. Logical forms are conceptual representations that are strictly determined by the grammar (and as such belong to the language system), but as conceptual representations they can also be elaborated or modified through non-linguistic considerations. The bare logical form of a sentence is a conceptual representation that is determined strictly by the grammar, while the modified logical form is a conceptual representation which has been shaped, in part, by extralinguistic factors such as world knowledge and contextual expectations.

The RT view seems to me intermediate between the two positions. 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 (the level of propositional forms) and, for that reason, they cannot be given a truth-theoretic interpretation.⁵

There is yet another position. So far I have assumed that there are two distinct systems : the language system and the conceptual system. According to some philosophers, however, 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. If that is so, then we can account for modified logical forms without appealing to a second system *in addition to* the language system. If a sentence is uttered and assigned, in context, a modified logical form resulting from the operation of free pragmatic processes, there are actually two sentences at play. What I have called the modified logical form is in fact the (bare) logical form of another sentence, that which runs through the speaker’s and/or the interpreter’s mind.

There are two possible versions of this view. One of them is rather familiar ; it can be found in the writings of theorists such as Jerry Katz and Kent Bach. The leading idea is this. Sentences have literal interpretations, but they can also be used to convey something different from (e.g. more determinate than) their strictly literal interpretation. When that happens,

⁵ 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 interpretively ambiguous with either quantifier capable of bearing wide scope » (Hornstein 1995 : 20). This is noticeable, for that implies that « May (1985) drops the requirement that sentences be disambiguated at LF » (*id.*). So there is a continuum of positions : if we define LF as that level of grammatical representation that incorporates whatever features of sentences structure enter directly into the semantic interpretation of sentences, the question arises as to how much pragmatic/conceptual elaboration is needed to get to *lfs*, the logical formulae which are input to the truth-theoretic machinery which delivers truth-conditions. One possible answer (favoured by researchers in the mainstream tradition) is: none – LF directly maps to *lf*, without any pragmatic/conceptual processing. As we have just seen, someone like May has to admit that some pragmatic/conceptual processing is needed in some cases e.g. to disambiguate the LF in order to get the pair-list reading for ‘What did everyone bring ?’. Relevance theorists (and TCP-theorists more generally) think a lot of pragmatic/conceptual processing is needed. (For more on ambiguous or underspecified logical forms, see van Deemter and Peters 1996.)

pragmatic processing maps the sentence they utter to some other sentence that was not uttered but might have been. Let us call the sentence that was actually uttered s_1 and the other sentence s_2 . The modified logical form of the utterance is the (bare) logical form of s_2 .⁶ This is how I understand Bach's notion of 'expansion' (Bach 1987, chapter 4 ; Bach 1994). On Bach's view the sentences one utters are often elliptical for more complex sentences one has in mind, in the nonlinguistic, sellarsian sense of 'elliptical'.⁷ For example, I can say 'There is a lion in the middle of the piazza' and mean that there is a *statue representing a lion* in the middle of the piazza. This contextual interpretation determines the utterance's modified logical form, and that is actually the logical form of a more complex sentence that was not uttered, but was running through my mind ('There is a statue representing a lion in the middle of the piazza'). On Bach's view, the free pragmatic process of 'expansion' is a linguistic process : it maps a sentence s_1 to another sentence s_2 by adding elements to s_1 .

The other possible interpretation of the view is less familiar, but it has recently found advocates in the linguistic and philosophical community. 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, 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.

Take the sentence 'There is a lion in the middle of the piazza' again, and assume that as a result of (what I take to be) an optional process of modulation, the word 'lion' here is understood as it is in the phrase 'stone lion'. Then, according to the view under discussion, what I take to be the modified logical form of the sentence *is* its logical form, and what I call its bare logical form is not its logical form at all: it is the logical form of the distinct,

⁶ As Katz puts it, « the utterance meaning of a sentence S can be expressed as the grammatical meaning of another sentence S' » (Katz 1977 : 19).

⁷ On the sellarsian sense of 'elliptical', see Neale 2000 : 286-287.

homophonous sentence ‘There is a lion in the middle of the piazza’ which means that there is a real lion in the middle of the piazza. What allegedly distinguishes the two sentences is the occurrence in the first one, but not in the second one, of a covert, optional element, e.g. a covert metonymic operator (or whatever accounts for the modulation of ‘lion’ in this context). The element in question has the following properties :

1. It is covert – that is why there is no superficial difference between the two sentences.
2. It is optional, hence it is always possible for what looks superficially like the same sentence not to carry that covert element and therefore not to have the meaning that results from the addition of that element.

Elements that have those properties I call ‘covert optionals’. By positing the existence of such elements in the language, one can account for the effects of free pragmatic processes while claiming that they are not pragmatic processes at all, but regular processes of semantic interpretation applied to covert elements. Such a view has been put forward by Luisa Marti (2006), by Josef Stern (2000, 2006), and by Polly Jacobson (2005).⁸

As an example, take metaphor, discussed by Stern. An expression is interpreted metaphorically, according to Stern, if and only if a covert ‘Mthat’-operator applies to that expression — a context-sensitive operator for which Stern supplies a Kaplan-inspired semantics. The ‘Mthat’-operator is optional : whenever it occurs, it is also possible to build a sentence indistinguishable from the metaphorical sentence but with a different meaning (since the alternative sentence does not carry the ‘Mthat’-operator that is responsible for the metaphorical interpretation). Stern suggests that the same sort of account will work for metonymy. Similarly, Luisa Marti, in her discussion of my views, posits covert optionals to account for all the cases for which I appeal to free enrichment. Whenever I invoke a free pragmatic function that makes the meaning of an expression more specific, Marti posits a covert variable *g* which is assigned, in context, that very function as its semantic interpretation. On that picture the alleged difference between free enrichment and saturation is simply a difference between two types of covert elements : those which, like the *g* variable in

⁸ One may also interpret in this light the ‘syntactic’ analysis of scalar implicatures put forward by Danny Fox (2005). On that analysis scalar implicatures result from the free insertion of a covert exhaustivity operator *Exh* with a meaning akin to that of ‘only’. See Fox and Hackl 2006 : 543.

question, are optional and can be omitted without making the sentence ungrammatical or otherwise deviant, and those which cannot be omitted. In ‘John is short’, a covert variable (for a comparison class or whatever serves as implicit parameter) is also involved but it is not optional : whenever what looks superficially like the sentence ‘John is short’ is uttered, the covert element has to be there. But the covert elements that account for metaphors, metonymies, free enrichment etc. are characterized by their optionality. They may be generated in the syntax but they need not be.⁹ As Marti writes,

The crucial difference between the two proposals resides in what bears the responsibility for optionality. In Recanati’s system, that is the responsibility of the pragmatics, of the properties of the context of utterance. In the system proposed here, the pragmatics has the same responsibility it has in the interpretation of pronouns, and only that. That is, given a variable in the syntax/at LF, there has to be a variable-assignment, which of course depends on the context of utterance, that provides values for this variable. But the pragmatics does not trigger anything in the sense of Recanati ; there is no process of free enrichment. Whether [a] variable is generated in the syntax or not is left completely free... The system tries out different derivations, and only those that comply with all the principles of grammar, including Gricean principles, are successful. (Marti 2006 : 149-150)

How are we to account for the difference between covert optionals and other covert elements lacking the optionality feature? Marti insists that there is only one sort of covert variable : the difference between the two types of case (alleged ‘saturation’ cases and alleged ‘modulation’ cases) is simply that something in the sentence imposes the presence of the covert variable in some cases (e.g. ‘short’ does) while in other cases the presence of the covert element is not imposed by anything in the sentence and could be omitted without ungrammaticality. In commenting on Marti’s paper, Polly Jacobson suggested that covert optionals are nothing but *covert adjuncts*. It is of the essence of adjuncts to be optional, she

⁹ When they are generated, they must be semantically interpreted and, if they are variables, they must be assigned a contextual value. What is optional is their generation, not their interpretation.

said, since their type is *a/a*.¹⁰ In contrast, standard saturation variables fill an argument place and cannot be omitted without ungrammaticality (Marti 2006 : 146-47).

Whatever we think of the line pursued by Marti, Jacobson and Stern, I think it has to be counted as another — admittedly deflationary — syntactic construal of free pragmatic processes. The main difference with the other two syntactic accounts is that *everything is now done within the language system* : on this account, what I call pragmatic modulation takes place through (i) the free generation of additional elements in the (covert) syntax, and (ii) the semantic interpretation of those elements along familiar lines. The resulting view sounds diametrically opposed to TCP, but the appearances may be deceptive. As far as I can tell, the only substantial difference there is between that account and the other syntactic accounts is that the level of syntactic representation to which the additional elements belong remains within the confines of the language system (rather than involving a shift to the conceptual system). What this difference exactly amounts to — what its consequences are — remains to be determined.

References

- Bach, Kent (1987) *Thought and Reference*. Oxford : Clarendon Press.
Bach, Kent (1994) ‘Conversational Implicature’, in *Mind and Language* 9 : 124-62.
Bach, Kent (2000) ‘Quantification, Qualification and Context’, in *Mind and Language* 15 : 262-83.

¹⁰ In a related vein, Jacobson shows that quantifier domain restriction (a phenomenon which Bach 2000 treats as an instance of ‘expansion’) can be accounted for in terms of *covert relative clauses* freely adjoined to the nouns (Jacobson 2005). Thus ‘every girl’ can be either :

every [_N **girl**] (contextually unrestricted reading)

or :

every [_N **girl** [_{RC} **PRO**]] (contextually restricted reading).

Contextual domain restriction on this account is a matter of assigning a value to a silent variable. As always with covert optionals, the assignment is obligatory ; what is optional is the generation of the variable.

- Carruthers, Peter (1996) *Language, Thought, and Consciousness*. Cambridge : Cambridge University Press.
- Carston, Robyn (2002) *Thoughts and Utterances : The Pragmatics of Explicit Communication*. Oxford : Blackwell.
- Chierchia, Gennaro (1999) 'Linguistics and Language', in Wilson, Robert and Keil, Frank (eds.) *The MIT Encyclopedia of the Cognitive Sciences*, Cambridge, Mass : MIT Press, xcix-cix.
- Chierchia, Gennaro (2004) 'Scalar Implicatures, Polarity Phenomena, and the Syntax/Pragmatics Interface', in Beletti, A. (ed.) *Structures and Beyond*, Oxford : Oxford University Press, 39-103.
- Chierchia, Gennaro and McConnell-Ginet, Sally (1990) *Meaning and Grammar*. Cambridge, Mass. : MIT Press/Bradford Books.
- Chomsky, Noam (1976) 'Conditions on Rules of Grammar', in *Linguistic Analysis* 2 : 303-51.
- van Deemter, Kees and Peters, Stanley (eds.) (1996) *Semantic Ambiguity and Underspecification*. Stanford : CSLI Publications.
- Fodor, Jerry (2001) 'Language, Thought and Compositionality', in *Mind and Language* 16 : 1-15.
- Fox, Danny (2005) 'Implicature Calculation, Pragmatics or Syntax ?' Class handout, Ecole Normale Supérieure.
- Fox, Danny and Hackl, Martin (2006) 'The Universal Density of Measurement', in *Linguistics and Philosophy* 29 : 537-86.
- Hornstein, Norbert (1995) *Logical Form : From GB to Minimalism*. Oxford : Blackwell.
- Jackendoff, Ray (1983) *Semantics and Cognition*. Cambridge, Mass : MIT Press.
- Jackendoff, Ray (1990) *Semantic Structures*. Cambridge, Mass : MIT Press.
- Jackendoff, Ray (1993) 'The Combinatorial Structure of Thought : The Family of Causative Concepts', in Reuland, Eric and Abraham, Werner (eds.) *Knowledge and Language II : Lexical and Conceptual Structure*, Dordrecht : Kluwer, 31-49.
- Jackendoff, Ray (1997) *The Architecture of the Language Faculty*. Cambridge, Mass. : MIT Press.
- Jackendoff, Ray (2002) *Foundations of Language*. Oxford : Oxford University Press.
- Jacobson, Polly (2005) 'Variable-Free Semantics : The Case of Quantifier Domain Restrictions'. Handout, Institut Jean-Nicod.
- Katz, Jerry (1972) *Semantic Theory*. New York : Harper and Row.
- Katz, Jerry (1977) *Propositional Structure and Illocutionary Force*. New York : Crowell.

- Kempson, Ruth (1993) 'Input Systems, Anaphora, Ellipsis and Operator Binding', in Reuland, Eric and Abraham, Werner (eds.) *Knowledge and Language II : Lexical and Conceptual Structure*, Dordrecht : Kluwer, 51-78.
- King, Jeffrey and Stanley, Jason (2005) 'Semantics, Pragmatics, and the Role of Semantic Content', in Szabo, Zoltan (ed.) *Semantics versus Pragmatics*, Oxford : Clarendon Press, 111-64.
- Ludlow, Peter (1999) *Semantics, Tense, and Time*. Cambridge, Mass : MIT Press.
- Marti, Luisa (2006) 'Unarticulated Constituents Revisited', in *Linguistics and Philosophy* 29 : 135-66.
- May, Robert (1985) *Logical Form*. Cambridge, Mass. : MIT Press.
- Neale, Stephen (2000) 'On being Explicit : Comments on Stanley and Szabo, and on Bach', in *Mind and Language* 15 : 284-94.
- Pagin, Peter (2005) 'Compositionality and Context', in Preyer, Gerhard and Peter, Georg (eds.) *Contextualism in Philosophy : Knowledge, Meaning, and Truth*, Oxford : Clarendon Press, 303-48.
- Predelli, Stefano (2005) *Contexts : Meaning, Truth, and the Use of Language*. Oxford : Clarendon Press.
- Recanati, François (2002) 'Unarticulated Constituents', in *Linguistics and Philosophy* 25 : 299-345.
- Recanati, François (2004) *Literal Meaning*. Cambridge : Cambridge University Press.
- Sperber, Dan and Wilson, Deirdre (1986) *Relevance : Communication and Cognition*. Oxford : Blackwell.
- Stanley, Jason (2000) 'Context and Logical Form', in *Linguistics and Philosophy* 23 : 391-434.
- Stanley, Jason (2005) 'Semantics in Context', in Preyer, Gerhard and Peter, Georg (eds.) *Contextualism in Philosophy : Knowledge, Meaning, and Truth*, Oxford : Clarendon Press, 221-53.
- Stern, Josef (2000) *Metaphor in Context*. Cambridge, Mass : MIT Press.
- Stern, Josef (2006) 'Metaphor, Literal, and Literalism', in *Mind and Language* 21 : 243-79.
- Szabo, Zoltan (2000) *Problems of Compositionality*. New York : Garland.