Introduction

Around the middle of the twentieth century, there were two opposing camps within the analytic philosophy of language. The first camp – IDEAL LANGUAGE PHILOSOPHY, as it was then called – was that of the pioneers, Frege, Russell, Carnap, Tarski, and so on. They were, first and foremost, logicians studying formal languages and, through them, 'language' in general. They were not originally concerned with natural language, which they thought defective in various ways;[1] yet, in the 1960s, some of their disciples established the relevance of their methods to the detailed study of natural language.[2] Their efforts gave rise to contemporary FORMAL SEMANTICS, a very active discipline whose stunning developments in the last quarter of the twentieth century changed the face of linguistics.

The other camp was that of so-called ORDINARY LANGUAGE PHILOSOPHERS, who thought important features of natural language were not revealed but hidden by the logical approach initiated by Frege and Russell. They advocated a more descriptive approach and emphasized the pragmatic nature of natural language as opposed to, say, the language of *Principia Mathematica*. Their own work[3] gave rise to contemporary pragmatics, a discipline which, like formal semantics, developed successfully within linguistics in the past forty years.

Central in the ideal language tradition had been the equation of, or at least the close connection between, the meaning of a (declarative) sentence and its truth-conditions. This truth-conditional approach to meaning is perpetuated, to a large extent, in contemporary formal semantics. A language is viewed as a system of rules or conventions, in virtue of which certain assemblages of

[1] There are a few exceptions. The most important one is Hans Reichenbach, whose insightful 'Analysis of conversational language' was published as a chapter – the longest – in his *Elements of Symbolic Logic* (Macmillan, 1947).

[2] See Richard Montague, *Formal Philosophy: Selected Papers* (Yale University Press, 1974), and Donald Davidson, *Inquiries into Truth and Interpretation* (Clarendon Press, 1984).

[3] The most influential authors were Austin, Strawson, Grice and the later Wittgenstein. Grice is a special case, for he had, as he once said, one foot in each of the two camps (Paul Grice, 'Retrospective Epilogue', in his *Studies in the Way of Words* (Harvard University Press, 1989), p. 372).

symbols count as well-formed, meaningful sentences. The meaning of a sentence (or of any complex symbol) is determined by the meanings of its parts and the way they are put together. Meaning itself is patterned after reference. The meaning of a simple symbol is the conventional assignment of a worldly entity to that symbol: for example, names are assigned objects, monadic predicates are assigned properties or sets of objects, and so on. The meaning of a declarative sentence, determined by the meanings of its constituents and the way they are put together, is equated with its truth-conditions. For example, the subject-predicate construction is associated with a semantic rule for determining the truth-conditions of a subject-predicate sentence on the basis of the meaning assigned to the subject and that assigned to the predicate. On this picture, knowing a language is like knowing a 'theory' by means of which one can deductively establish the truth-conditions of any sentence of that language.

This truth-conditional approach to meaning is one of the things which ordinary language philosophers found quite unpalatable. According to them, reference and truth cannot be ascribed to linguistic expressions in abstraction from their use. In vacuo, words do not refer and sentences do not have truth-conditions. Words–world relations are established through, and indissociable from, the use of language. It is therefore misleading to construe the meaning of a word as some worldly entity that it represents or, more generally, as its truth-conditional contribution. The meaning of a word, insofar as there is such a thing, should rather be equated with its use-potential or its use-conditions. In any case, what must be studied primarily is speech: the activity of saying things. Then we will be in a position to understand language, the instrument we use in speech. Austin's theory of speech acts and Grice's theory of speaker's meaning were both meant to provide the foundation for a theory of language, or at least for a theory of linguistic meaning.

Despite the early antagonism I have just described, semantics (the formal study of meaning and truth-conditions) and pragmatics (the study of language in use) are now conceived of as complementary disciplines, shedding light on different aspects of language. The heated arguments between ideal language philosophers and ordinary language philosophers are almost forgotten. There are two main reasons for the new situation. On the one hand semanticists, in moving from artificial to natural languages, have given up Carnap's idea that the semantic relation between words and the world can be studied in abstraction from the context of use.[4] That the Carnapian abstraction is illegitimate given the pervasiveness of context-sensitivity in natural language is fully acknowledged by those working in formal semantics. On the other hand those

[4] See my 'Pragmatics and Semantics', in Larry Horn and Gregory Ward (eds.), *Handbook of Pragmatics* (Blackwell, forthcoming).

working in pragmatics no longer hold that 'meaning is use'. Instructed by Grice, they systematically draw a distinction between what a given expression means, and what *its use* means or conveys, in a particular context (or even in general).

Still, the ongoing debate about the best delimitation of the respective territories of semantics and pragmatics betrays the persistence of two recognizable currents or approaches within contemporary theorizing. According to the dominant position, which I call 'Literalism', we may legitimately ascribe truth-conditional content to natural language *sentences*, quite independently of what the speaker who utters this sentence means. 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.

I say that Literalism is the dominant position because I believe most philosophers of language and linguists would accept the following description of the division of labour between semantics and pragmatics:

> Semantics deals with the literal meaning of words and sentences as determined by the rules of the language, while pragmatics deals with what users of the language mean by their utterances of words or sentences. To determine 'what the speaker means' is to answer questions such as: Was John's utterance intended as a piece of advice or as a threat? By saying that it was late, did Mary mean that I should have left earlier? Notions such as that of illocutionary force (Austin) and conversational implicature (Grice) thus turn out to be the central pragmatic notions. In contrast, the central semantic notions turn out to be reference and truth. It is in terms of *these* notions that one can make explicit what the conventional significance of most words and expressions consists in.
>
> The meaning of an expression may be insufficient to determine its referential content: that is so whenever the expression is indexical or otherwise context-dependent. In such cases, the meaning of the expression provides a rule which, given a context, enables the interpreter to determine the content of the expression in that context. The content thus determined in context by the conventional meanings of words is their *literal content*. The literal content of a complete declarative utterance is 'what is said', or the proposition expressed, by that utterance.
>
> As Grice emphasized, a speaker's meaning is not a matter of rules but a matter of intentions: what someone means is what he or she overtly intends (or, as Grice says, 'M-intends') to get across through his or her utterance. Communication succeeds when the M-intentions of the speaker are recognized by the hearer. Part of the evidence used by the hearer in working out what the speaker means is provided by the literal content of the uttered sentence, to which the hearer has independent access via his knowledge of the language. In ideal cases of linguistic communication, the speaker means exactly what she says, and no more is required to understand the speech act than a correct understanding of the sentence uttered in performing it. In real life, however, what the speaker means typically goes beyond, or otherwise diverges from, what the uttered sentence literally says. In such cases the hearer must rely on background knowledge to determine what the speaker means – what her communicative intentions are.

There is much that is correct in this description, but there also is something which I think must be rejected, namely the *contrast* between literal truth-conditions and speaker's meaning. That contrast commits us to Literalism, and in this book I want to argue for Contextualism. According to Contextualism, the contrast between what the speaker means and what she literally says is illusory, and the notion of 'what the sentence says' incoherent. What is said (the truth-conditional content of the utterance) is nothing but an aspect of speaker's meaning. That is not to deny that there *is* a legitimate contrast to be drawn between what the speaker says and what he or she merely implies. Both, however, belong to the realm of 'speaker's meaning' and are pragmatic through and through.

I will not only criticize Literalism and argue for Contextualism in the following chapters. I will discuss all sorts of intermediate positions corresponding to views actually held in the current debate about the semantics/pragmatics interface. Whether or not one accepts my arguments, I hope the survey of logical space which I provide will be useful to those interested in the debate, and will contribute to shaping it in the years to come.

1 Two approaches to 'what is said'

1.1 The basic triad

Anyone who has reflected on the sentence meaning/speaker's meaning distinction knows that a simple distinction is in fact insufficient. Two equally important distinctions must be made. First, there is the distinction between the linguistic meaning of a sentence-type, and what is said (the proposition expressed) by an utterance of the sentence. For example, the English sentence 'I am French' has a certain meaning which, *qua* meaning of a sentence-type, is not affected by changes in the context of utterance. This context-independent meaning contrasts with the context-dependent propositions which the sentence expresses with respect to particular contexts. Thus 'I am French', said by me, expresses the proposition that I am French; if you utter the sentence, it expresses a different proposition, even though its linguistic meaning remains the same across contexts of use.

Second, there is a no less important distinction between what is actually said and what is merely 'conveyed' by the utterance. My utterance of 'I am French' expresses the proposition that I am French, but there are contexts in which it conveys much more. Suppose that, having been asked whether I can cook, I reply: 'I am French.' Clearly my utterance (in this context) provides an affirmative answer to the question. The meaning of the utterance in such a case includes more than what is literally said; it also includes what the utterance 'implicates'.[1]

'What is said' being a term common to both distinctions, we end up with a triad:

sentence meaning
vs
what is said
vs
what is implicated

[1] See Paul Grice, *Studies in the Way of Words* (Harvard University Press, 1989), p. 24: 'I wish to introduce, as terms of art, the verb *implicate* and the related nouns *implicature* (cf. *implying*) and *implicatum* (cf. *what is implied*). The point of this manoeuvre is to avoid having, on each occasion, to choose beween this or that member of the family of verbs for which *implicate* is to do general duty.'

The distinguishing characteristic of sentence meaning (the linguistic meaning of the sentence type) is that it is conventional and context-independent. Moreover, in general at least, it falls short of constituting a complete proposition, that is, something truth-evaluable. In contrast, both 'what is said' and 'what is implicated' are context-dependent and propositional. The difference between 'what is said' and 'what is implicated' is that the former is constrained by sentence meaning in a way in which the implicatures aren't. What is said results from fleshing out the meaning of the sentence (which is like a semantic 'skeleton') so as to make it propositional. The propositions one can arrive at through this process of contextual enrichment or 'fleshing out' are constrained by the skeleton which serves as input to the process. Thus 'I am French' can express an indefinite number of propositions, but the propositions in question all have to be compatible with the semantic potential of the sentence; this is why the English sentence 'I am French' cannot express the proposition that kangaroos have tails. There is no such constraint on the propositions which an utterance of the sentence can communicate through the mechanism of implicature. Given enough background, an utterance of 'I am French' might implicate that kangaroos have tails. What's implicated is implicated by virtue of an inference, and the inference chain can (in principle) be as long and involve as many background assumptions as one wishes.

The basic triad can be mapped back onto the simple sentence meaning/speaker's meaning distinction by grouping together two of the three levels. There are two ways to do it, corresponding to two interpretations for the triad. The 'minimalist' interpretation stresses the close connection between sentence meaning and what is said; together, sentence meaning and what is said constitute the *literal meaning* of the utterance as opposed to what *the speaker* means:

literal meaning { sentence meaning / what is said
vs
speaker's meaning

The other, 'non-minimalist' interpretation of the triad stresses the commonality between what is said and what is implicated, both of which are taken to be pragmatically determined:

sentence meaning
vs
speaker's meaning { what is said / what is implicated

Essential to this interpretation is the claim that 'what is said', though constrained by the meaning of the sentence, is not as tightly constrained as is traditionally thought and, in particular, does not obey what I will refer to as the 'minimalist' constraint.

1.2 Minimalism

As I said above, what distinguishes 'what is said' from the implicatures is the fact that the former must be 'closely related to the conventional meaning of the words (the sentence) [one] has uttered'.[2] However, this constraint can be construed more or less strictly. What I call 'Minimalism' construes the constraint very strictly; 'what is said', in the minimalist framework, departs from the conventional meaning of the sentence (and incorporates contextual elements) *only when this is necessary to 'complete' the meaning of the sentence and make it propositional*. In other words, the distance between sentence meaning and what is said is kept to a minimum (hence the name 'Minimalism').

The crucial notion here is that of 'saturation'. Saturation is the process whereby the meaning of the sentence is completed and made propositional through the contextual assignment of semantic values to the constituents of the sentence whose interpretation is context-dependent (and, possibly, through the contextual provision of 'unarticulated' propositional constituents, if one assumes, as some philosophers do, that such constituents are sometimes needed to make the sentence fully propositional). This process takes place whenever the meaning of the sentence includes something like a 'slot' requiring completion or a 'free variable' requiring contextual instantiation.[3] Thus an indexical sentence like 'He is tall' does not express a complete proposition unless a referent has been contextually assigned to the demonstrative pronoun 'he', which acts like a free variable in need of contextual instantiation. Genitives provide another well-known example: an utterance including the phrase 'John's book' does not express a complete proposition unless a particular relation has been identified as holding between the book and John. Nominal compounds work the same way: 'burglar nightmare' means something like 'a nightmare that bears a certain relation *R* to burglars', which relation must be contextually identified. Other well-known examples of saturation include parametric predicates ('small', 'on the left'), definite null instantiation (that is, the case where one of the arguments in the semantic structure of a lexeme, typically a verb, is not syntactically realized and must be contextually identified, as when someone says 'I heard' or 'I noticed'), and so on and so forth.

Whenever saturation is in order, appeal to the context is necessary for the utterance to express a complete proposition: from a semantic point of view, saturation is a *mandatory* contextual process. Other contextual processes – for example, the inference process generating implicatures – are semantically

[2] Grice, *Way of Words*, p. 25.

[3] Even when saturation consists in contextually providing a constituent that is unarticulated in surface syntax (as the implicit argument in 'I noticed'), it is something in the sentence (here the predicate 'notice', which arguably denotes a two-place relation) which triggers the search for the contextual element and makes it obligatory. See §2.1 of my 'Unarticulated Constituents', in *Linguistics and Philosophy* 25 (2002), 299–345.

optional in the sense that the aspects of meaning they generate are dispensable; the utterance would still express a complete proposition without them. According to Minimalism, those extra constituents of meaning which are not necessary for propositionality are external to what is said. The only justification for including some pragmatically determined constituent of meaning into what is said (as opposed to what is merely conveyed) is the indispensability of such a constituent – the fact that the utterance would not express a complete proposition if the context did not provide such a constituent.

1.3 Literal truth-conditions vs actual truth-conditions

Consider examples (1)–(6), often discussed in the literature:

(1) I've had breakfast.
(2) You are not going to die.
(3) It's raining.
(4) The table is covered with books.
(5) Everybody went to Paris.
(6) John has three children.

In all such cases, as we shall see, the minimalist constraint implies that what the utterance literally says is not what intuitively seems to be said.

From a minimalist point of view, the first sentence, 'I've had breakfast', expresses the proposition that S (the speaker) has had breakfast before t* (the time of utterance). Strictly speaking this proposition would be true if the speaker had had breakfast twenty years ago and never since. This is clearly not what the speaker means (when she answers the question 'Do you want something to eat?' and replies 'I've had breakfast'); she means something much more specific, namely that she's had breakfast *on that very day* (that is, the day which includes t*). This aspect of speaker's meaning, however, has to be construed as external to what is said and as being merely conveyed, in the same way in which the utterer of 'I am French' implies, but does not say, that he is a good cook. That is so because the 'minimal' interpretation, to the effect that the speaker's life was not entirely breakfastless, is sufficient to make the utterance propositional. Nothing in the sentence itself forces us to bring in the implicit reference to a particular time span. Indeed we can easily imagine contexts in which a speaker would use the same sentence to assert the minimal proposition and nothing more.[4]

The same thing holds even more clearly for the second example. Kent Bach, to whom it is due, imagines a child crying because of a minor cut and her mother uttering (2) in response. What is meant is: 'You're not going to die from that cut.' But literally the utterance expresses the proposition that the kid will not die *tout court* – as if he or she were immortal. The extra element contextually

[4] Dan Sperber and Deirdre Wilson, *Relevance: Communication and Cognition* (Blackwell, 1986), pp. 189–90. For an alternative analysis of that example, see my 'Pragmatics of What is Said', in *Mind and Language* 4 (1989), pp. 305–6, and §6.2 below.

provided (the implicit reference to the cut) does not correspond to anything in the sentence itself; nor is it an unarticulated constituent whose contextual provision is necessary to make the utterance fully propositional. Again, we can easily imagine a context in which the same sentence would be used to communicate the minimal proposition and nothing more.[5]

What about (3)? John Perry and many others after him have argued as follows.[6] Even though nothing in the sentence 'It's raining' stands for a place, nevertheless it does not express a complete proposition unless a place is contextually provided. The verb 'to rain', Perry says, denotes a dyadic relation – a relation between times and places. In a given place, it doesn't just rain or not, it rains at some times while not raining at others; similarly, at a given time, it rains in some places while not raining in others. To evaluate a statement of rain as true or false, Perry says, we need both a time and a place. Since the statement 'It is raining' explicitly gives us only the two-place relation (supplied by the verb) and the temporal argument (indexically supplied by the present tense), the relevant locational argument must be contextually supplied for the utterance to express a complete proposition. If Perry is right, the contextual provision of the place concerned by the rain is an instance of saturation, like the assignment of a contextual value to the present tense: both the place and the time are constituents of what is said, even though, unlike the time, the place remains unarticulated in surface syntax.

But is Perry right? If really the contextual provision of a place was mandatory, hence an instance of saturation, *every* token of 'It's raining' would be unevaluable unless a place were contextually specified. Yet I have no difficulty imagining a counterexample, that is, a context in which 'It is raining' is evaluable even though no particular place is contextually singled out. In 'Unarticulated Constituents' I depicted an imaginary situation in which

> rain has become extremely rare and important, and rain detectors have been disposed all over the territory (whatever the territory – possibly the whole Earth). In the imagined scenario, each detector triggers an alarm bell in the Monitoring Room when it detects rain. There is a single bell; the location of the triggering detector is indicated by a light on a board in the Monitoring Room. After weeks of total drought, the bell eventually rings in the Monitoring Room. Hearing it, the weatherman on duty in the adjacent room shouts: 'It's raining!' His utterance is true, iff it is raining (at the time of utterance) in some place or other.[7]

The fact that one can imagine an utterance of 'It's raining' that is true iff it is raining (at the time of utterance) in some place or other arguably establishes

[5] Kent Bach, 'Conversational Impliciture', in *Mind and Language* 9 (1994), p. 134. For an alternative analysis of that example (in terms of domain restriction), see below § 6.2.

[6] John Perry, 'Thought Without Representation' (1986), reprinted (with a postscript) in his collection *The Problem of the Essential Indexical and Other Essays* (Oxford University Press, 1993), 205–25.

[7] Recanati, 'Unarticulated Constituents', p. 317.

the pragmatic nature of the felt necessity to single out a particular place, in the contexts in which such a necessity is indeed felt. When a particular place is contextually provided as relevant to the evaluation of the utterance, it is for pragmatic reasons, not because it is linguistically required. (Again, if it were linguistically required, in virtue of semantic properties of the sentence type, it would be required in *every* context.) If this is right, then the contextual provision of a place is not an instance of saturation after all: it's not something that's mandatory. It follows (by minimalist standards) that the place is not a constituent of what is strictly and literally said: when I say 'It is raining' (rather than something more specific like 'It's raining in Paris' or 'It's raining here'), what I *literally* say is true iff it's raining somewhere or other.[8] That is obviously not what I mean, since what I mean involves a particular place. Appearances notwithstanding, the situation is similar to the case of 'I've had breakfast', where a restricted time interval is contextually provided for pragmatic reasons, without being linguistically mandated.

Examples (4) and (5) are amenable to the same sort of treatment. According to standard Russellian analysis, a definite description conveys an implication of uniqueness: hence 'The table is covered with books' is true iff there is one and only one table and it is covered with books. To make sense of this, we need either to focus on a restricted situation in which there is indeed a single table, or to expand the predicate 'table' and enrich it into, say, 'table of the living-room' in order to satisfy the uniqueness constraint. Either way, it is arguable that the form of enrichment through which we make sense of the utterance is not linguistically mandated: it is only pragmatically required. If we don't enrich, what we get is an already complete proposition (albeit one that is pretty absurd): the proposition that the only existing table is covered with books. Similarly with example (5): without enrichment the utterance expresses a proposition that is true iff every existing person went to Paris. Such a proposition is unlikely to be true, but that does not make it incomplete. On this view the enrichment process through which, in context, we reach the proposition actually communicated (to the effect that everybody *in such and such group* went to Paris) is not linguistically but pragmatically required; hence it is not an instance of saturation, but an optional process of 'free enrichment'. It follows that, in those examples as much as in the previous ones, the proposition literally expressed is different from, and more general than, the proposition actually communicated.

1.4 A problem for Minimalism

In general, the literal truth-conditions posited as part of the minimalist analysis turn out to be very different from the intuitive truth-conditions which

[8] See Emma Borg, 'Saying What You Mean: Unarticulated Constituents and Communication' (forthcoming) for a defence of that claim.