

Literal Meaning

François Recanati

Institut Jean-Nicod, Paris



PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS

The Edinburgh Building, Cambridge, CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa

<http://www.cambridge.org>

© François Recanati 2004

This book is in copyright. Subject to statutory exception and to the provisions of relevant collective licensing agreements, no reproduction of any part may take place without the written permission of Cambridge University Press.

First published 2004

Printed in the United Kingdom at the University Press, Cambridge

Typeface Times 10/12 pt. *System* L^AT_EX 2_ε [T_B]

A catalogue record for this book is available from the British Library

Library of Congress cataloguing-in-publication data

Recanati, François, 1952–

Literal meaning / François Recanati.

p. cm.

Includes bibliographical references and index.

ISBN 0 521 79246 0 (hardback) – ISBN 0 521 53736 3 (paperback)

1. Semantics. 2. Semantics (Philosophy). 3. Pragmatics. I. Title.

P325.R344 2003 401'.43 – dc21 2003051475

ISBN 0 521 79246 0 hardback

ISBN 0 521 53736 3 paperback

Contents

<i>Acknowledgments</i>	<i>page vii</i>
Introduction	1
1 Two approaches to ‘what is said’	5
1.1 The basic triad	5
1.2 Minimalism	7
1.3 Literal truth-conditions vs actual truth-conditions	8
1.4 A problem for Minimalism	10
1.5 The availability of what is said	13
1.6 The availability based approach	16
1.7 ‘Saying’ as a pragmatic notion	18
1.8 Availability vs Minimalism	20
2 Primary pragmatic processes	23
2.1 Enrichment, loosening and transfer	23
2.2 Rejecting the Gricean picture	27
2.3 Accessibility	30
2.4 Objections and responses	32
2.5 Interactive processing	34
2.6 The role of schemata	36
3 Relevance-theoretic objections	38
3.1 One or two systems?	38
3.2 Personal and sub-personal inferences	40
3.3 Implicature or enrichment?	44
3.4 Mutual adjustment of explicature and implicature	46
3.5 Implicated premisses	48
3.6 Personal-level inferences: occurrent vs dispositional	49
4 The Syncretic View	51
4.1 Four levels?	51
4.2 Semantics and pragmatics: the literalist picture	54
4.3 Semantic underdeterminacy	56
4.4 The minimal proposition as ‘common denominator’	58
4.5 Interaction between saturation and optional pragmatic processes	61
4.6 Do we need the minimal proposition?	64
4.7 The reflexive proposition	65

5	Non-literal uses	68
5.1	Non-literal uses as non-minimal departures from literal meaning	68
5.2	Non-literal uses and secondary meaning	70
5.3	Non-minimal departures without secondariness	72
5.4	The transparency condition	74
5.5	Varieties of non-literal meaning	75
5.6	Internal vs external duality	78
5.7	Conclusion	81
6	From Literalism to Contextualism	83
6.1	Five positions	83
6.2	Indexicalism	86
6.3	Contextualism	90
6.4	Literalist responses to the contextualist challenge	92
6.5	Where Indexicalism and Contextualism meet	95
7	Indexicalism and the Binding Fallacy	98
7.1	Mandatory vs optional	98
7.2	Two criteria	100
7.3	The indexicalist challenge	103
7.4	Is the Binding Criterion reliable?	105
7.5	Variadic functions	107
7.6	The Binding Fallacy	109
7.7	Conclusion: the failure of Indexicalism	111
8	Circumstances of evaluation	115
8.1	Modality	115
8.2	Time and tense	118
8.3	Situations	121
8.4	Saturation or enrichment?	124
8.5	Sub-sentential circumstances	125
8.6	Conclusion	127
9	Contextualism: how far can we go?	131
9.1	The modulation of sense	131
9.2	The semantic relevance of modulation	133
9.3	Four approaches	136
9.4	Truth-conditional unstability: from Waismann's 'open texture' to Searle's 'background'	141
9.5	Ostensive definitions	144
9.6	Meaning Eliminativism	146
9.7	Conclusion	151
	Conclusion	154
10.1	Alleged arguments against Contextualism	154
10.2	Remnants of Literalism	159
10.3	Availability, Minimalism, and the dispositional/occurrent contrast	162
	<i>Bibliography</i>	166
	<i>Index</i>	175

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 between 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 Implicature', 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.

untutored conversational participants would ascribe to the utterance. This divergence between the intuitive truth-conditions of an utterance and the literal truth-conditions postulated by the theorist is particularly striking in connection with examples like (6). According to a fairly standard view,⁹ the proposition literally expressed by (6) is the proposition that John has at least three children, that is, no less than three but possibly more. In certain contexts this corresponds to what the speaker actually means (as when I say, 'If John has three children he can benefit from lower rates on public transport') but in other contexts what the speaker means is quite different. Suppose for example that I am asked how many children John has and that I reply by uttering (6). Clearly, in this context, I mean that John has (exactly) three children – no more and no less. This is standardly accounted for by saying that the proposition literally expressed, to the effect that John has at least three children, combines with the 'implicature' that John has no more than three children (a generalized implicature that is accounted for in terms of the maxim of quantity);¹⁰ as a result of this combination, what is globally communicated – and what I actually mean – is the proposition that John has exactly three children. Now *this is the only proposition I am conscious of expressing by my utterance*; in particular, I am unaware of having expressed the 'minimal' proposition that John has at least three children. To account for this obvious fact, the minimalist claims that we are aware only of what is globally conveyed or 'communicated' by the utterance. Analysing this into 'what is literally said' and 'what is implied' is the linguist's task, not something that is incumbent upon the normal language user. Figure 1.1 (p. 12) illustrates this widespread conception.

The problem with this conception is that it lacks generality. Recall the example I gave earlier – the utterance 'I am French' used to convey that I am a good cook. In the relevant situation of utterance, both the speaker and the listener are aware that the speaker says he is French, and thereby implies he is a good cook. This typical case of implicature is very different from a case like (6) in which the speaker is not only (like the hearer) unaware of the proposition literally expressed, but would strongly deny having said what the minimalist claims was actually said.

It turns out that there are two sorts of case. On the one hand there are prototypical cases of implied meaning, in which the participants in the speech situation are aware both of what is said and of what is implied, and also of the inferential

⁹ See Larry Horn, *The Natural History of Negation* (Chicago University Press, 1989), pp. 205–16.

¹⁰ As Grice puts it in one of his early papers, 'one should not make a weaker statement rather than a stronger one unless there is a good reason for so doing' (Paul Grice, 'The Causal Theory of Perception', in *Proceedings of the Aristotelian Society*, Supplementary Volume 35 (1961), p. 132). Since the statement that John has (at least) three children is weaker than the statement that John has n children (for $n > 3$), the maxim is obeyed only if John has no more than three children. (If John has more than three children, the statement that he has three is too weak and violates the maxim.) The statement 'John has three children' therefore implicates that John has no more than three children, in virtue of the presumption that the maxim is obeyed.

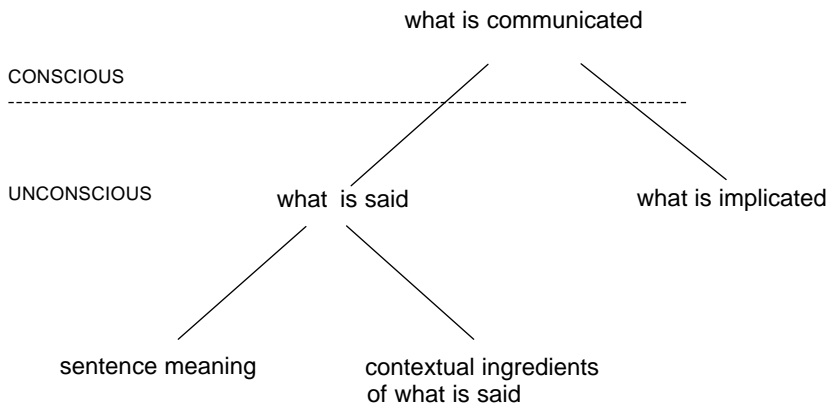


Figure 1.1 The standard approach

connection between them. On the other hand, there are the cases illustrated by (1)–(6). Given his willingness to treat certain aspects of the intuitive meaning of (1)–(6) as conversational implicatures external to what is literally said, the minimalist must explain why those implicatures, unlike the prototypical cases (for instance the French/cook example), do not have the property of conscious ‘availability’.

The only explanation I have come across in the literature makes use of Grice’s distinction between ‘generalized’ and ‘particularized’ conversational implicatures, that is, between implicatures which arise ‘by default’, without any particular context or special scenario being necessary, and those which require such specific contexts. In contrast with the latter, the former are ‘hard to distinguish from the *semantic* content of linguistic expressions, because such implicatures [are] routinely associated with linguistic expressions in all ordinary contexts’.¹¹ Generalized implicatures are unconsciously and automatically generated and interpreted. They belong to the ‘micropragmatic’ rather than to the ‘macropragmatic’ level, in Robin Campbell’s typology:

A macropragmatic process is one constituted by a sequence of explicit inferences governed by principles of rational cooperation. A micropragmatic process develops as a cryptic [= unconscious] and heuristic procedure which partially replaces some macropragmatic process and which defaults to it in the event of breakdown.¹²

¹¹ Stephen Levinson, *Pragmatics* (Cambridge University Press, 1983), p. 127. See also his *Presumptive Meanings: The Theory of Generalized Conversational Implicature* (MIT Press, 2000).

¹² Robin Campbell, ‘Language Acquisition, Psychological Dualism and the Definition of Pragmatics’, in Herman Parret, Marina Sbisà and Jeff Verschueren (eds.), *Possibilities and Limitations of Pragmatics* (Benjamins, 1981), p. 101.

But there are problems with this explanation. According to Horn, the generalized nature of an implicature does not entail its conscious unavailability – its 'cryptic' character.¹³ In other words, it is possible for an implicature to be both 'generalized' and intuitively accessible as an implicature distinct from what is said. Thus Horn insists that the generalized scalar implicature from 'some' to 'not all' *is* consciously available (in contrast to that from 'three' to 'exactly three'). A speaker saying 'Some students came to the meeting' normally implies that not all students came, and when this is so there is no tendency on the part of the interpreter to conflate the implicature with what is said. This is actually debatable, for the 'implicature' at issue can arise at sub-sentential level (for example, 'He believes some students came'), and in such cases there are reasons to doubt that the availability condition is satisfied. Be that as it may, the 'generalization' of an implicature does not seem to be necessary for its unconscious character. Many particularized 'bridging' inferences are automatic and unconscious. To take an example from Robyn Carston, 'He went to the cliff and jumped' is readily interpreted as saying that the person referred to jumped *over the cliff*, even though this is only contextually suggested.

1.5 The availability of what is said

In earlier writings I put forward a conception diametrically opposed to that illustrated by figure 1.1 above.¹⁴ 'What is said', I held, *is* consciously available to the participants in the speech situation (figure 1.2). 'What is communicated' is not a distinct level where 'what is said' and 'what is implied' have been merged and integrated into a unified whole; it is merely a name for the level at which we find both what is said and what is implied, which level is characterized by conscious accessibility. On this picture, there are only two basic levels: the bottom level at which we find both the meaning of the sentence and the contextual factors which combine with it to yield what is said; and the top level at which we find both what is said and what is implied, both being consciously accessible (and accessible as distinct).

The availability of what is said follows from Grice's idea that saying itself is a variety of non-natural meaning. One of the distinguishing characteristics of non-natural meaning, on Grice's analysis, is its essential overttness. Non-natural meaning works by openly letting the addressee recognize one's primary intention (for example, the intention to impart a certain piece of information, or the intention to have the addressee behave in a certain way), that is, by

¹³ Larry Horn, 'The Said and the Unsaid', in Chris Barker and David Dowty (eds.), *SALT 2: Proceedings of the Second Conference on Semantics and Linguistic Theory* (Ohio State University Working Papers in Linguistics 40, 1992), 163–92.

¹⁴ See 'The Pragmatics of What is Said', already cited; *Direct Reference: From Language to Thought* (Blackwell, 1993), pp. 233–74; and 'What is Said', in *Synthese* 128 (2001), 75–91.

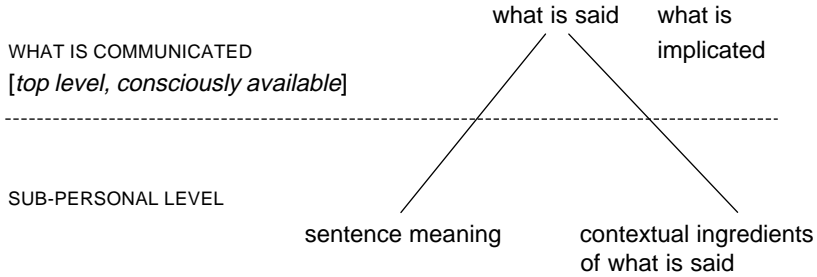


Figure 1.2 An alternative approach

(openly) expressing that intention so as to make it graspable. This can be done in all sorts of ways, verbal or non-verbal. Even if we restrict ourselves to verbal communication, there are many ways in which we can mean things by uttering words. *Saying* is one way; *implying* is another.

The view that 'saying' is a variety of non-natural meaning entails that what is said (like what is meant in general, including what is implied) *must* be available – it must be open to public view. That is so because non-natural meaning is essentially a matter of intention-recognition. On this view what is said by uttering a sentence depends upon, and can hardly be severed from, the speaker's publicly recognizable intentions. Hence my 'Availability Principle', according to which 'what is said' must be analysed in conformity to the intuitions shared by those who fully understand the utterance¹⁵ – typically the speaker and the hearer, in a normal conversational setting.

I take the conversational participants' intuitions concerning what is said to be revealed by their views concerning the utterance's truth-conditions. I assume that whoever fully understands a declarative utterance knows which state of affairs would possibly constitute a truth-maker for that utterance, that is, knows in what sort of circumstance it would be true. The ability to pair an utterance with a type of situation in this way is more basic than, and in any case does not presuppose, the ability to *report* what is said by using indirect speech; it does not even presuppose mastery of the notion of 'saying'. Thus the proper way to elicit such intuitions is not to ask the subjects 'What do you think is said (as opposed to implied or whatever) by this sentence as uttered in that situation'?¹⁶

¹⁵ Recanati, *Direct Reference*, p. 248.

¹⁶ Michael Thau notes that: 'speakers almost never explicitly think about the distinction between what they've said and what they've implicated. So the question of what a speaker takes himself to have said by some utterance will have to depend upon the answer he *would* give if he *were* asked. And it's very likely that in many circumstances there won't be a single answer, that the answer will differ depending on how the question is put. It's also very likely that the answer will vary from circumstance to circumstance' (*Consciousness and Cognition* (Oxford University Press, 2002), p. 148). Contrary to what Thau thinks, however, this does not speak against the

I therefore tend to agree with Bach’s criticism of the experiments through which Gibbs and Moise attempted to support the availability based approach:¹⁷

[They] thought they could get their data about what is said, and thereby test the validity of Recanati’s Availability Principle, by asking people what is said by a given utterance, or by asking them whether something that is conveyed by a given utterance is implicated or merely said. Evidently they assume that what people *say* about what is said is strongly indicative of what *is* said. In fact, what it is indicative of is how people apply the phrase ‘what is said’ . . . It tells us little about what is said, much less about the cognitive processes whereby people understand utterances.¹⁸

However, Bach himself uses what he calls the ‘IQ test’ to determine what is said, that is, *he ties what is said to indirect speech reports of what is said*.¹⁹ I find this procedure most objectionable, and that is *not* what I mean when I claim that what is said should be individuated according to the intuitions of normal interpreters. Thus I strongly disagree with Cappelen and Lepore’s surprising statement:

We ourselves don’t see how to elicit intuitions about what-is-said by an utterance of a sentence without appealing to intuitions about the accuracy of indirect reports of the form ‘He said that . . .’ or ‘What he said is that . . .’ or even ‘What was said is that . . .’²⁰

I find this statement surprising, because there obviously *is* another way of eliciting truth-conditional intuitions. One has simply to provide subjects with scenarios describing situations, or, even better, with – possibly animated – pictures of situations, and to ask them to evaluate the target utterance as true or false with respect to the situations in question.²¹ That procedure has been used by several researchers to test speaker’s intuitions about, for example, the truth-conditions of donkey sentences. Thus Bart Geurts describes his experimental set-up (inspired from earlier work by Yoon) as follows:

Twenty native speakers of Dutch were asked to judge whether or not donkey sentences correctly described pictured situations. Instructions urged subjects to answer either true

availability based approach. The speaker’s intuitions concerning what is said need *not* involve the very notion of what is said.

¹⁷ Raymond Gibbs and Jessica Moise, ‘Pragmatics in Understanding What is Said’, in *Cognition* 62 (1997), 51–74.

¹⁸ Kent Bach, ‘Seemingly Semantic Intuitions’, in Joseph Keim Campbell, Michael O’Rourke and David Schier (eds.), *Meaning and Truth* (Seven Bridges Press, 2002), p. 27.

¹⁹ ‘IQ’ means INDIRECT QUOTATION. On the IQ test, see Bach’s papers ‘Semantic Slack’, in Savas Tsohatzidis (ed.), *Foundations of Speech Act Theory* (Routledge, 1994), 267–91, ‘The Myth of Conventional Implicature’, in *Linguistics and Philosophy* 22 (1999), 327–66, and ‘You Don’t Say?’, in *Synthese* 128 (2001), 15–44.

²⁰ Herman Cappelen and Ernie Lepore, ‘On an Alleged Connection Between Indirect Speech and the Theory of Meaning’, in *Mind and Language* 12 (1997), p. 280.

²¹ For an implicit use of that procedure, see Saul Kripke, *Naming and Necessity* (Blackwell, 1980), p. 12 .

or false, but they were also given the option of leaving the matter open in case they couldn't make up their minds.²²

This procedure presupposes that normal interpreters have intuitions concerning the truth-conditional content of utterances. On my view, those intuitions correspond to a certain 'level' in the comprehension process – a level that a proper theory of language understanding must capture. That is the level of 'what is said' (as opposed to, for example, what is implied).

1.6 The availability based approach

From a psychological point of view, we can draw a helpful parallel between understanding what one is told and understanding what one sees. In vision, the retinal stimuli undergo a complex (multi-stage) train of processing which ultimately outputs a conscious perception, with the dual character noted by Brentano: the subject is aware both of what he sees, and of the fact that he is seeing it. Although more complex in certain respects, the situation with language is similar. The auditory signal undergoes a multi-stage train of processing which ultimately outputs a conceptual experience: the subject understands what is said. This is very much like (high-level) perception. If I am told that it is four o'clock, I hear that it is four o'clock, just as, when I look at my watch, I see that it is four o'clock. Like the visual experience, the locutionary experience possesses a dual character: we are aware both of what is said, and of the fact that the speaker is saying it.

In calling understanding an *experience*, like perception, I want to stress its conscious character. Understanding what is said involves entertaining a mental representation of the subject-matter of the utterance that is both determinate enough (truth-evaluable) *and* consciously available to the subject. This suggests a criterion, distinct from the minimalist criterion, for demarcating what is said. Instead of looking at things from the linguistic side and equating 'what is said' with the minimal proposition one arrives at through saturation, we can take a more psychological stance and equate what is said with (the semantic content of) the conscious output of the complex train of processing which underlies comprehension.²³

To be sure, that output itself is subject to further processing through, for example, inferential exploitation. Consider, once again, vision. Seeing John's car, I can infer that he is around. Similarly, hearing that John has had breakfast, I can infer that he is not hungry and does not need to be fed. Just as what

²² Bart Geurts, 'Donkey Business', in *Linguistics and Philosophy* 25 (2002), p. 135.

²³ As Ian Rumfitt once put it, 'what is said in the course of an utterance is nothing other than what somebody who understands the utterance understands to be said' ('Content and Context: the Paratactic Theory Revisited and Revised', in *Mind* 102 (1993), p. 439).

is seen corresponds to the primary conscious output of visual processing, not to what can be secondarily derived from it, 'what is said' corresponds to the primary truth-evaluable representation made available to the subject (at the personal level)²⁴ as a result of processing the sentence. It is therefore minimal in a certain sense, though not (as we shall see) in the sense of Minimalism.

Accordingly, I distinguish between two sorts of pragmatic process. The contextual processes which, like saturation, are (sub-personally) involved in the determination of what is said I call *primary* pragmatic processes. In contrast, *secondary* pragmatic processes are ordinary inferential processes taking us from what is said, or rather from the speaker's saying of what is said, to something that (under standard assumptions of rationality and cooperativeness) follows from the fact that the speaker has said what she has said. To the extent that the speaker overtly intends the hearer to recognize such consequences as following from her speech act, they form an integral part of what the speaker means by her utterance. That is, roughly, Grice's theory of 'conversational implicature'. An essential aspect of that theory is that the hearer must be able to recognize what is said and to work out the inferential connection between what is said and what is implied by saying it. Again, it follows that what is said must be consciously available to the interpreter. It must satisfy what I call the availability constraint.

In this framework we solve the difficulty raised in section 1.5. We no longer have two sorts of case of implicature – the prototypical cases where the interlocutors are aware of what is said, aware of what is implied, and aware of the inferential connection between them, and the cases in which there is no such awareness. Conscious awareness is now a built-in feature of both what is said and the implicatures. That is so because what is said is the conscious output of linguistic-cum-pragmatic processing, and the implicatures correspond to further conscious representations inferentially derived, at the personal rather than sub-personal level, from what is said (or, rather, from the speaker's saying what is said). The alleged cases in which the speech participants themselves are not distinctly aware of what is said and of what is implied are reclassified: they are no longer treated as cases of 'implicature', strictly speaking, but as cases in which a primary pragmatic process operates in the (sub-personal) determination of what is said.²⁵

²⁴ On the contrast between the personal and sub-personal levels, see Daniel Dennett, *Content and Consciousness* (Routledge and Kegan Paul, 1969), pp. 93–6, and 'Toward a Cognitive Theory of Consciousness', in his *Brainstorms: Philosophical Essays on Mind and Psychology* (MIT Press, 1981), p. 153.

²⁵ This is consonant with the approach taken by some semanticists who insist that, for example, scalar 'implicatures' 'are not computed *after* truth-conditions of (root) sentences have been figured out; they are computed phrase by phrase' (Gennaro Chierchia, 'Scalar Implicatures, Polarity Phenomena, and the Syntax/Pragmatics Interface', forthcoming). In chapter 2, I will stress the fact that primary pragmatic processes operate locally, in contrast to secondary pragmatic processes, which can only operate when the truth-conditions of the sentence have been worked out.

1.7 ‘Saying’ as a pragmatic notion

So far I have followed Grice, who construes saying as a variety of meaning. But this pragmatic approach to ‘saying’ is controversial. Most philosophers use the notion of ‘what is said’ (or ‘the proposition expressed’) in such a way that it is *not* a ‘pragmatic’ notion – having to do with what the speaker means or with what the hearer understands. What is said is supposed to be a property of the *sentence* (with respect to the context at hand) – a property which it has in virtue of the rules of the language.

Minimalism is closely associated with such a non-pragmatic way of looking at what is said. In the minimalist framework, saturation is the only contextual process allowed to affect ‘what is said’, because it alone is a *bottom-up* process, that is, a process triggered (and made obligatory) by a linguistic expression in the sentence itself.²⁶ All other contextual processes determine aspects of meaning external and additional to what is said. Take, for example, ‘free enrichment’ – the process responsible for making the interpretation of an utterance more specific than its literal interpretation (as when ‘jumped’ is contextually understood as ‘jumped over the cliff’). That form of enrichment is ‘free’ in the sense of not being linguistically controlled. Thus what triggers the contextual provision of the relevant temporal restriction in example (1) (‘I’ve had breakfast’) is not something in the sentence but simply the fact that the utterance is meant as an answer to a question about the speaker’s present state of hunger (which state can be causally affected only by a breakfast taken on the same day). While saturation is a bottom-up, linguistically controlled pragmatic process, free enrichment is a top-down, pragmatically controlled pragmatic process. Insofar as it is pragmatically rather than linguistically controlled, free enrichment is taken to be irrelevant to ‘what is said’, on the non-pragmatic construal of what is said.

I will discuss the non-pragmatic construal of what is said in chapter 4. For the time being, I’m interested in the pragmatic construal, based on Grice’s idea, and the reasons it provides for rejecting the minimalist constraint (§1.8). Before turning to that issue, however, I want to rebut a couple of objections to the pragmatic construal.

The first objection is this. If, following Grice, we construe saying as a variety of meaning, we will be prevented from acknowledging an important class of cases in which the speaker does not mean what he says. Irony is a good example of that class of cases. If I say ‘John is a fine friend’ ironically, in a context in which it is obvious to everybody that I think just the opposite, it is clear that I do not mean what I say: I mean the opposite. Still, I *say* that John is a fine friend. Grice’s construal of saying as a variety of meaning prevents him from

²⁶ As I pointed out in footnote 3, p. 7, that is true even when saturation consists in providing a so-called ‘unarticulated constituent’.

acknowledging that fact. According to Grice, when I say 'John is a fine friend' in the mentioned situation, I do not *really* say that John is a fine friend – I *pretend* to be saying it. The pragmatic construal of saying forces Grice to draw a distinction between 'saying' and 'making as if to say'.

As far as I am concerned, I find Grice's distinction (between genuine saying and making as if to say) perfectly legitimate, but I can understand the worries of those who feel that the notion of 'saying' he uses is too much on the pragmatic, illocutionary side.²⁷ We certainly need a notion of 'what is said' which captures *the objective content of an utterance irrespective of its pragmatic force as a serious assertion or as an ironical utterance*. Still, I find the objection superficial, for it is quite easy actually to construct the desired notion within Grice's own framework. Grice uses 'say' in a strict sense. In that sense whatever is said must be meant. But we can easily define a broader sense for 'say':

S says that *p*, in the broad sense, iff he either says that *p* (in the strict sense) or makes as if to say that *p* (again, in the strict sense of 'say').

I will henceforth use 'say' in that broad sense, which remains within the confines of the pragmatic construal.

Another objection to the pragmatic construal focuses on the loss of objectivity that allegedly goes with it. What is said is objective in the sense that it is possible both for the speaker to make a mistake and say something other than what he means, and for the hearer to misunderstand what the speaker is saying. Those mistakes are possible, the objector will argue, because what is said is an objective property of the sentence (in context). But on the pragmatic construal, it is not clear that this objectivity can be captured. Imagine the following situation: the speaker wants to say that Paul is tall, and, mistaking Tim for Paul, says 'He is tall' while pointing to Tim. The speaker thus inadvertently says that Tim is tall. Now imagine that the hearer also mistakes Tim for Paul. Thanks to this lucky mistake, he grasps what the speaker means, thinking that this is what he has said. The speaker and the hearer therefore converge on a certain interpretation, which is not objectively what was said, but which they both (mistakenly) think is what was said. How, in the framework I have sketched, will it be possible to dissociate what is actually said from the protagonists' mistaken apprehension of what is said? Have we not equated what is said with their understanding of what is said?

We have not. We have equated what is said with what a *normal interpreter* would understand as being said, in the context at hand. A normal interpreter knows which sentence was uttered, knows the meaning of that sentence, knows

²⁷ 'The verb "say", as Grice uses it, does not mark a (locutionary) level distinct from that marked by such illocutionary verbs as "state" and "tell", but rather functions as a generic illocutionary verb' (Bach, 'You Don't Say?', p. 41).

the relevant contextual facts (who is being pointed to, and so on).²⁸ Ordinary users of the language *are* normal interpreters, in most situations. They know the relevant facts and have the relevant abilities. But there are situations (as in the above example) where the actual users make mistakes and are not normal interpreters. In such situations their interpretations do not fix what is said. To determine what is said, we need to look at the interpretation that a normal interpreter would give. This is objective enough, yet remains within the confines of the pragmatic construal.

1.8 Availability vs Minimalism

In the framework I have sketched, there is a basic constraint on what is said:

Availability

What is said must be intuitively accessible to the conversational participants (unless something goes wrong and they do not count as ‘normal interpreters’).

This constraint leads us to give up Minimalism. That is the price to pay if we want Availability to be satisfied.

The reason why Availability is incompatible with Minimalism is simple enough. The aspects of the meaning of (1)–(6) which the minimalist construes as conversational implicatures are, one may admit, contextual ingredients in the overall meaning of the utterance. They do *not* belong to the conventional meaning of the sentence. The minimalist claims that they do not belong to ‘what is said’ either, because they are optional: those contextual aspects of the meaning of the utterance are not necessary for the latter to express a complete proposition. But the availability constraint pulls in the other direction. The very fact that the minimal propositions allegedly expressed are not consciously available shows that it would be a mistake to equate them to what is said; rather, the availability constraint dictates that the aspects of meaning which Minimalism construes as external to what is said (for example, the implicit reference to a place in (3), or to the cut in (2), or to a time interval in (1)) are actually *constitutive* of what is said, because *when we subtract them from the intuitive meaning of the utterance the proposition which results is no longer something accessible to the participants in the speech situation*. Thus we have two quite distinct phenomena: examples like ‘I am French’/‘I am a good cook’ involve something which is said and whose saying implies something else; examples like (1)–(6), in contrast, do not involve the distinction between what is said and what is implied but a different distinction between the literal meaning of the sentence and contextual ingredients entering into the determination of what

²⁸ This is all tacit knowledge, not the sort of ‘conscious awareness’ I talk about in connection with secondary pragmatic processes.

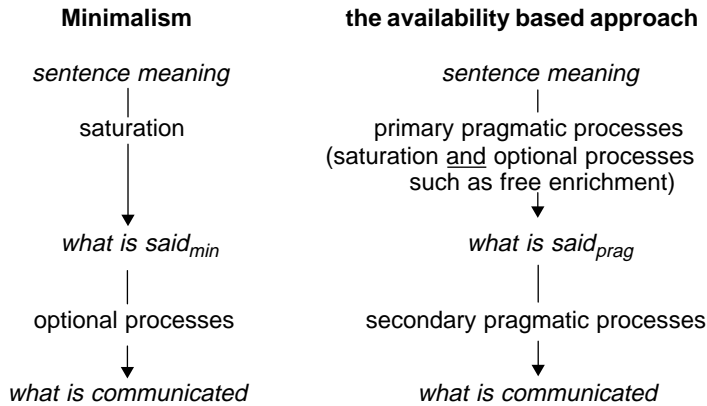


Figure 1.3 Comparing the approaches

is said. If we maintain that those ingredients are indeed ‘optional’ rather than necessary for propositionality, this implies that *we must give up the minimalist criterion according to which the context contributes to what is said only when this is necessary for some proposition to be expressed.*

According to the view we arrive at, truth-conditional interpretation is pragmatic to a large extent. Various pragmatic processes come into play in the very determination of what is said; not merely saturation – the contextual assignment of values to indexicals and free variables in the logical form of the utterance – but also free enrichment and other processes which are not linguistically triggered but are pragmatic through and through. Figure 1.3 summarizes the contrast between the two conceptions (Minimalism, and the availability based approach).

According to the availability based approach, the crucial distinction is not between mandatory and optional contextual processes, but between those that are ‘primary’ and those that are ‘secondary’. Primary pragmatic processes include not only saturation, but also ‘optional’ processes such as free enrichment. Independent evidence for their inclusion in this category is provided by the fact that, in general, the notion of ‘what is said’ we need to capture the input to secondary, inferential processes already incorporates contextual elements of the optional variety. Consider examples (1)–(6) once again. In each case we may suppose that the speaker implies various things by saying what she does. Thus, by saying that she’s had breakfast, the speaker implies that she is not hungry and does not want to be fed. By saying that the child is not going to die, the mother implies that the cut is not serious; and so forth. Now those implicatures can be worked out only if the speaker is recognized as expressing the (non-minimal) proposition that she’s had breakfast *that morning*, or that the child won’t die *from that cut*. Clearly, if the speaker had had breakfast twenty

years ago (rather than that very morning), nothing would follow concerning the speaker's present state of hunger and her willingness or unwillingness to eat something. The implicature could not be derived, if what the speaker says was not given the richer, temporally restricted interpretation. If therefore we accept the Gricean picture, according to which 'what is said' serves as input to the secondary process of implicature-generation, we must, *pace* Grice himself, acknowledge the non-minimal character of what is said. This provides some support to the availability based approach, as against Minimalism.