

Contextualism, Metaphor, and What is Said

ELISABETH CAMP

Abstract: On a familiar and *prima facie* plausible view of metaphor, speakers who speak metaphorically *say* one thing in order to *mean* another. A variety of theorists have recently challenged this view; they offer criteria for distinguishing what is said from what is merely meant, and argue that these support classifying metaphor within ‘what is said’. I consider four such criteria, and argue that when properly understood, they support the traditional classification instead. I conclude by sketching how we might extract a workable notion of ‘what is said’ from ordinary intuitions about saying.

1. Contextualism and ‘What is Said’

Metaphor is a deeply context-sensitive linguistic phenomenon. In the right context, nearly any term or sentence can be used metaphorically, and can be used to express a wide variety of contents. The standard way to accommodate this broad variability is to treat metaphor as a form of speaker meaning, on which speakers intentionally say one thing in order to communicate something different (Grice, 1975; Searle, 1979; Martinich, 1984). On this view, by uttering:

- (1) Bill is a mouse

Alice knowingly says something which, *if* she meant it, would commit her to the claim that Bill is a small rodent. Normally, she won’t intend to be taken as committing herself to such an absurdly false claim; her hearers realize this, and interpret her metaphorically instead. The particular assumptions she intends her hearers to employ in determining her metaphorical content can vary considerably across different conversational contexts, producing a wide variety of possible metaphorical meanings.

This broadly Gricean model is both intuitively plausible and theoretically satisfying. It nicely subsumes metaphor within a larger theory of communication, on which speakers intentionally exploit shared conversational presuppositions in order to communicate efficiently. And it accomplishes this while allowing us to retain an attractive view of the relation between speaker and sentence meaning. On Grice’s own preferred way of thinking, ‘what is said’ by an utterance of a

Thanks to Jay Atlas, Kent Bach, Anne Bezuidenhout, Anthony Corsentino, Andy Egan, Catherine Elgin, John Hawthorne, Jeff King, John MacFarlane, Bernard Nickel, Marga Reimer, Dmitri Tymoczko, and Catherine Wearing. Thanks also to two anonymous referees for *Mind & Language*, and to audiences at talks at the Pacific APA and Rutgers University from which this paper is drawn.

Address for correspondence: Society of Fellows, Harvard University, 78 Mount Auburn Street, Cambridge MA 02138, U.S.A.

Email: ecamp@fas.harvard.edu

sentence *S* is closely related to ‘the particular meanings of the elements of *S*, their order, and their syntactic character’ (Grice, 1989, p. 87). This close connection between conventionally encoded sentence meaning and ‘what a speaker says’ forges a straightforward connection between the semantic and pragmatic components of linguistic theory, while preserving a robust form of compositionality. This in turn is desirable, because compositionality offers our best hope for explaining the massive generativity of our linguistic capacities.

Despite these benefits, a growing number of philosophers and linguists find this overall picture of communication to be a pretty but hopelessly oversimplified philosopher’s fantasy. The linguistic reality, they think, is that context-specific assumptions and expectations interact with established linguistic conventions in a much more intimate and pervasive fashion. We often—even typically—use words to express thoughts that depart from their conventionally encoded meaning. On the traditional understanding of ‘what is said’, it follows that speakers very often speak indirectly and intentionally say false things. This can seem inherently implausible. More importantly, it can also appear methodologically suspect. If the traditional ‘Gricean’ is allowed to treat the uses to which we ordinarily put our words in such a cavalier fashion, then it looks like he can postulate any semantics he pleases. He can always insist that the speaker *said* just what his semantics predicts, and invoke his elaborate pragmatic machinery to bridge the gap between that and what is intuitively communicated.¹

‘Contextualists’ argue that rather than postulating such elaborate, psychologically unmotivated machinery, we should allow that contextual factors can pervasively ‘intrude’ into the semantics, to affect what a speaker says itself. ‘What is said’ is first and foremost a feature of utterances, they think. It is therefore inappropriate simply to stipulate that ‘what is said’ by an utterance is the uttered sentence’s conventionally encoded semantic meaning, as many advocates of the traditional picture (though not Grice himself) do. Instead, contextualists think, ‘what is said’ should be defined in terms of what speakers *do* in uttering their words. More specifically, they argue, saying involves the direct and explicit expression of one’s intended meaning. Speakers can express their thoughts directly and explicitly in some conversational contexts in ways they could not in other contexts, because different particular background assumptions happen to be in play in those contexts.

The model cases for this view are ‘meaning enrichment’ and loose talk. An utterance of:

- (2) Have you had lunch?

is a case of enrichment: it would typically be understood as asking whether the hearer has eaten lunch *today*, and not just at any time in the past. In loose talk,

¹ Cf. Levinson (2000, p. 231): ‘If this tactic is pursued willy-nilly, in violation of our intuitions about truth and falsity, why not claim that any other sentence for which the proponent’s semantic theory makes the wrong predictions is in fact patched up by the postsemantic pragmatics and thus is after all correctly analyzed by his unlikely theory?’

satisfaction-conditions are stripped away from, rather than added to, conventionally encoded meaning. For instance, an utterance of:

- (3) It is silent in the house.

is naturally treated as true just in case it is really quiet, despite the distant drone of a refrigerator, drips from a leaky faucet, and the like, even though the lexical entry for 'silent' arguably requires a total absence of sound (Bezuidenhout, 2001b, p. 168). Both of these cases stand in sharp contrast to Grice's paradigmatic examples of implicature. If a professor writes a letter of recommendation consisting entirely of the sentence:

- (4) John's command of English is excellent, and his attendance at tutorials has been regular,

then intuitively, she hasn't explicitly stated her intended meaning—that John is a bad philosopher. All she has said is that his command of English is excellent and his attendance at tutorials regular; the burden falls on the hearer to work out what her actual, intended meaning might be. Contextualists think we should respect the important intuitive difference between these two classes of cases.

Several contextualists have recently argued that metaphor involves the same sort of temporary lifting and shifting of constraints on conventional meaning as loose talk and enrichment, so that words used metaphorically systematically also assume the role of expressing '*ad hoc*' concepts (Carston, 2002). They have offered four main arguments to support this claim that metaphor should be included within 'what is said'. First, ordinary speakers are normally willing to use 'say' to *report* metaphorically expressed contents. Second, metaphorical comprehension is *direct*, in the sense of coming first in the order of interpretation. Third, metaphorical speech can itself serve as a *vehicle* for sarcasm and implicature, but not vice versa. And finally, metaphor makes the speaker's intended content *explicit*, in the sense that hearers can respond to the speaker's intended content by echoing her words.

My aim in this paper is to defend the traditional view of metaphor against these arguments. In speaking metaphorically, I believe, speakers say one thing, which they typically don't mean, in order to mean something else.² I will argue that this view is

² There are cases in which speakers intend their utterance to be interpreted both literally and metaphorically; Ted Cohen (1976, p. 254) offers as an example:

- (i) Jesus was a carpenter.

David Hills (1997, p. 130) calls these metaphors 'twice apt': the speaker here does mean the sentence's literal meaning, but means something else in addition. Such cases contrast with merely 'twice true' metaphors, like

- (ii) No man is an island.

actually supported by the criteria the contextualists themselves offer, once those criteria are properly understood. Further, although my focus here is on metaphor specifically, I think the considerations I raise suggest a way to address the deeper contextualist objection that the traditional understanding of ‘what is said’ is an ungrounded theoretical fantasy. To this end, I conclude by sketching how we might extract a workable theoretical notion of ‘what is said’ from ordinary intuitions and practice.

Before we proceed, an important note of clarification. ‘Contextualists’ hold a variety of views about the relation between semantic content and ‘what is said’, which can make them difficult to discuss as a group. The traditional view identifies four things: conventionally encoded meaning, semantic meaning, ‘what is said’, and what is asserted (Figure 1). Some contextualists, such as Robyn Carston (1999; 2002), François Recanati (1995; 2004), and Anne Bezuidenhout (2001a; 2001b), retain the traditional assumption that semantic meaning is conventionally encoded meaning, but allow ‘what is said’ to come apart from semantic meaning. Others, such as John Perry (1986), Mark Crimmins (1998) and David Hills (1997), retain the traditional tight connection between semantics and ‘what is said’ while abandoning the conventionality of semantic meaning. (Some but not all of those who subscribe to this second view maintain that the contextual effects on word use are so pervasive that no substantive notion of ‘conventional word meaning’ can be sustained. On this more radical form of the view (cf. Davidson, 1986), there are only more or less overlapping idiolects. On the less radical form (cf. Clark, 1983), contextual effects operate on conventional meaning to alter semantic content.)




Figure 1 *The Traditional ‘Gricean’ View.*

In this case the literal meaning is true, but trivially so, and hence not part of what the speaker intends to communicate. ‘Twice apt’ metaphors might themselves seem to provide an apt demonstration of why metaphor in general should be treated as a form of speaker meaning rather than saying it’s natural to distinguish ‘twice apt’ metaphors from other metaphors on the ground that in the former case, the speaker means what she says and something else in addition, while in the typical case she doesn’t mean what she says. The contextualist must treat ‘twice apt’ metaphors as cases in which the speaker says both the literal sentence meaning and something else. This is possible, but sounds implausible to my ear. In any case, because ‘twice apt’ metaphors are unusual, I focus my attention on the typical case.

Whether or not they accept that semantic meaning is ‘what is said’, though, contextualists all agree that contextual factors pervasively ‘intrude’ into ‘what is said’; in this sense they all reject the traditional Gricean understanding of that phrase (Figure 2).

2. ‘What is Said’ as What is Ordinarily Said to be Said

The first and most direct argument for including metaphor within ‘what is said’ is that language users are often willing to report speakers who spoke metaphorically as having ‘said’ their intended contents. As an example, Bezuidenhout (2001b, p. 157) offers the following dialogue:

- (5) A: How about Bill?
- B: Bill’s a bulldozer.
- A: That’s true. But isn’t that a good thing in this case? We want someone who’ll stand up to the administration and get things done for our department.
- C: I disagree that he’s a bulldozer; that exterior hides someone who’s basically insecure. But either way, Bill wouldn’t make a good chair.

She writes that ‘the most natural construal to put on such a dialogue is that B says something, that B says it with assertoric force, and that A and C are either agreeing or disagreeing with what B says’; she might have added, more specifically, that it would be natural to report B as having *said* that Bill is a tough guy who doesn’t let obstacles stand in his way. Contextualists argue that the traditional notion of ‘what is said’ is systematically out of kilter with our ordinary use of ‘say’, and that we should redraw the boundary to reflect our ordinary ways of talking. Once we do this, metaphor too will fall within ‘what is said.’

It is undeniable that the ordinary use of ‘say’ and related terms is much more encompassing than the traditional theoretical sense on which ‘what is said’ is closely tied to conventional semantic meaning. The problem with redrawing the




Figure 2 *The Contextualist View.*

boundary of 'what is said' to reflect ordinary use, though, is that this use is so broad as to mark a different distinction altogether. For instance, the most normal ways to report utterances of:

- (6) I could use some salt on this pork chop.
- (7) I was wondering whether maybe you'd be willing to consider the possibility of writing a letter of recommendation for me.

respectively, are by saying something like:

- (6') She said she'd like the salt.
- (7') She asked for a letter of recommendation.

But (6) and (7) are paradigmatic examples of indirect speech. In each case, the speaker did mean the semantic content of her uttered sentence, but she was primarily concerned to make a further speech act; and an appropriate report will make this further intention explicit. Herman Cappelen and Ernie Lepore (1997, p. 285) use the following example to make a similar point. Suppose a professor is asked whether Alice passed her exam, and responds by saying:

- (8) I didn't fail any students.

Then (8') would be a natural report of the professor's utterance:

- (8') The professor said that Alice passed her exam.

But everyone agrees that the content attributed in (8') should count as a mere implicature of (8).

These are all examples of speakers meaning something *more* than the content of their uttered sentence's conventional meaning. It is often also acceptable to use 'say' to report the speaker's intended content when she means something *different* from her uttered sentence's conventionally encoded content. Thus, one might report a sarcastic utterance of:

- (9) What a brilliant idea: let's spend our last dollar on beer! Then I suppose we can hitchhike home.

with something like:

- (9') She's saying that we should save some money for the cab.

In this case, too, the gap between the utterance's conventional sentence meaning and the reported content is too great for even contextualists to include the latter within 'what was said'.

Cappelen and Lepore conclude from examples like this that that our ordinary use of 'say' is hopelessly undisciplined and theoretically useless. I don't think this conclusion follows: I believe that our ordinary practice of speech-reporting is sensitive to a certain standard of explicitness, and that speakers can legitimately object to reports that disregard this standard. Specifically, it is normally only appropriate to report speakers as having 'said' contents to which they have *openly and obviously* committed themselves by their utterance. For something to count as an implicature at all, the speaker must intend that her hearer 'work out' that meaning, and to do so because he recognizes this intention of hers. But within this constraint, there is considerable variation in how obvious the speaker intends her intention to be. And it is normally not appropriate to report speakers as having said things that they have merely insinuated or suggested, and so for which their communicative intentions are less than obvious. For instance, if I respond to a question about how good a philosopher John is by saying:

- (4) John's command of English is excellent, and his attendance at tutorials has been regular,

or to a question about how trustworthy he is by saying:

- (10) I think he's been out of trouble for quite a while now,

then I should not be reported as having said that John is untrustworthy or not a good philosopher. I have carefully chosen my answer so as to refrain from obviously committing myself to those claims. I could therefore legitimately object to such paraphrases by saying 'I didn't say *that!*'. By contrast, such an objection would normally ring quite hollow for any of the utterances (6) through (9) above.³

In illustrating the distinction between saying and merely meaning, contextualists have treated implicatures like (4) and (10) as paradigmatic. In fact, though, insinuating implicatures represent only a small part of the broader category of implicature and speaker meaning more generally, as examples (6) through (9) remind us. Once we take account of the full variety of implicature, it becomes clear that implicatures fall on both sides of the line between what speakers are and are not ordinarily willing to report people as having said.

Indeed, metaphor itself falls on both sides of the line between what it is and is not appropriate to report someone as saying. It is often *not* appropriate to report the speaker of a highly poetic metaphor as having said her metaphorical content. For instance, having considered:

- (11) The hourglass whispers to the lion's paw,

³ As we'll see below, there is a stricter use of 'say' on which the speaker didn't *say* in (6) the content reported in (6'). The speaker herself cannot appeal to this use of 'say' here, though, precisely because her intention to be interpreted as in (6') is so obvious.

in its original context of utterance (by W. H. Auden in 'Our Bias', translating a poem by Alex Sitnitsky), I am willing to hazard that the poet means something like that every source of activity and forcefulness is ultimately undone by the passage of time. I'm not at all certain of this. But even if I were, I'd still be very hesitant to report the poet as having *said* this, precisely because his metaphor is so clearly intended to be elusive and allusive.

3. 'What is Said' as 'First' or 'Primary' Meaning

The second argument for including metaphor within 'what is said' treats 'what is said' as meaning which comes first in the order of interpretation; the contextualist argues that metaphor is 'direct' or 'primary' in this sense. Thus, Bezuidenhout (2001b, p. 160) claims that the metaphorical interpretation of:

(12) He is a butcher.

'is not accessed by first considering and rejecting the literal interpretation of that sentence': interpretation proceeds straight to the metaphorical meaning. Similarly, Recanati (2001, p. 271) claims that an ordinary hearer 'unreflectively constructs' the metaphorical meaning of:

(13) The ATM swallowed my credit card,

'without going through a two-step procedure involving the prior computation of the 'literal' meaning (whatever that may be) and a secondary inference to the actual meaning'. By contrast, he argues, classic conversational implicatures like:

(14) Paul looks thirsty,

meaning that Paul might like a drink, and sarcastic utterances, like:

(15) Paul is a fine friend,

meaning that he is anything but, do require such a 'two-step computational procedure'.⁴

⁴ Bezuidenhout (p.c.) allows that irony, and perhaps even implicature, may not require a two-step comprehension process, because both 'explicatures' and implicatures can be derived in parallel on the basis of contextual clues. However, she does think that irony and implicature, but not metaphor, are 'indirect' in the sense articulated by Recanati's 'transparency condition' below.

The crucial question is how to understand the notion of ‘directness’ here, so that metaphorical meaning counts as direct while implicature and sarcasm do not. One possible construal appeals to sub-personal processing. As Bezuidenhout says, ‘Many reading time experiments have found that non-literal interpretations are accessed as quickly as literal ones’ (2001b, pp. 166–7). In fact, recent empirical evidence suggests that novel metaphors do take significantly longer to process (cf. Blasko and Connine, 1993; Brisard *et al.*, 2001; Bowdle and Gentner, 2005). In general, comprehension time (along with conscious processing) is affected more by the utterance’s immediate conversational relevance and by the explicit availability of relevant interpretive assumptions than by the utterance’s status as metaphorical, sarcastic, or indirect (Giora, 1997, 2002, 2003). Thus, sub-personal processing as reflected by comprehension times is not a good guide to whether metaphor should count as ‘what is said’.

Instead, both Recanati and Bezuidenhout ground ‘directness’ in ordinary language users’ intuitions about meaning; Recanati encapsulates this as a ‘transparency condition’ on indirectness, or what he calls ‘p-nonliterality’:⁵

P-nonliteralness is transparent to the language users ... This transparency is not a contingent property of p-nonliteralness. It is definitive of p-nonliteralness that the sort of inference at issue is conducted at the ‘personal’, rather than sub-personal, level and is therefore available to the language users (2001, pp. 270–1).

The relevant question thus becomes how to understand ‘transparency’. In this passage Recanati seems to assume that it requires language users to be consciously aware of going through an inferential process of deriving the secondary meaning. Understood in this sense, though, the transparency requirement is untenably strong: I’m not conscious of going through a two-stage inferential process to determine the relevant interpretation of utterances which contextualists do want to count as indirect, such as (14) and (15).⁶ In context, I just know what the speaker is getting at. And as Recanati admits (2001, p. 270), repeated use of a specific form of indirection, or explicit mention of contextual assumptions, will ‘short-circuit or conventionalize’ what would otherwise be a two-step interpretive process into a single stage—‘a complication,’ he says, ‘that I shall ignore’.

A weaker construal is that, as Recanati puts it above, the nonliteralness must be ‘available’ to ordinary speakers, in the sense that they must be aware that

⁵ Recanati stipulates that when language users arrive *directly* at an intended interpretation, and treat it as the *primary* meaning of the sentence, then the utterance is ‘p-literal’ (with ‘p’ standing for ‘primary’). To say that an utterance is ‘p-nonliteral,’ then, is just to say that it is indirect or ‘secondary’ in the order of interpretation.

⁶ Cf. García-Carpintero (2001).

there is something ‘special’ about the use of words. This is captured in another of his formulations: that ‘language users [must be] aware of the distinction between the two layers of meaning as well as of the connection between them’.⁷ This seems right: I believe that conversational implicatures and sarcastic utterances are nonliteral because I, like non-theorist speakers, recognize that my intuitive sense of their intended meanings differs markedly from, while depending upon, their conventional meanings. Further, this interpretation of the transparency requirement is nicely grounded in a plausible view of human reasoning and action in general. We often act intentionally and automatically, for reasons that are ‘available’ to us but which we don’t articulate explicitly even to ourselves. Linguistic interpretation is just a special case of this.⁸ As Kent Bach writes,

Communicative reasoning, like default reasoning in general, is a case of jumping to conclusions without consciously taking into account all alternatives or all relevant considerations. Even so, to be warranted such reasoning must be sensitive to such considerations. This means that such considerations can play a dispositional role even when they do not play an explicit role. They lurk in the background, so to speak, waiting to be taken into account when worth considering (2001, p. 259).

The trouble for the contextualist, though, is that the transparency requirement so construed *supports* the traditional analysis of metaphor. Ordinary speakers do think that metaphorical speech exploits established conventional meanings to novel ends, and so they do recognize that it ‘derives from a more basic, primary meaning which it presupposes’ (Recanati, 2001, p. 270). Further, when they are challenged to justify their interpretations of metaphorical utterances, ordinary language users do feel compelled to, and are in fact able to, articulate ‘the distinction between the two layers of meaning’ and then construct a rough rational reconstruction of ‘the connection between them.’ For example, they can and do offer justifications along the following rough lines:

Well, she can’t really be claiming that Bill is a *bulldozer*, because that’d be absurd: he’s a man, not a piece of landscape machinery. But he does share with bulldozers a propensity to obliterate obstacles. Since we’re talking about whether Bill would be a good department chair, she must mean that he would stand up to the administration and get things done.

⁷ This quote immediately precedes the strong formulation; Recanati intends the latter as a restatement of the former.

⁸ Cf. Lewis (1983, 181): ‘An action may be rational, and may be explained by the agent’s beliefs and desires, even though ... the agent gave no thought to the beliefs and desires which were his reason for acting.’

Their ability to articulate such patterns of reasoning reveals that they are indeed sensitive, at the 'personal level', to both the literal and intended meanings of metaphorical utterances. So literal meaning, and not metaphorical meaning, is 'direct' or 'first' meaning—by the contextualists' own criterion. The argument therefore fails to establish that metaphor belongs within 'what is said'.

Dead and dying metaphors like:

- (13) The ATM swallowed my credit card,
- (16) He vented his anger,

provide the strongest cases for a 'direct expression' view, because we do jump so easily and unreflectively to their metaphorical interpretations. However, it's important to remember that highly routinized metaphors like these lie at one extreme on a spectrum from conventionality to novelty. At the other extreme, poetic metaphors like:

- (11) The hourglass whispers to the lion's paw,

flaunt their literal absurdity, and require (and invite) sustained interpretive effort. In between, many ordinary conversational metaphors, such as:

- (17) Bill is a bulldozer,
- (18) My job is a jail,

are on a par with most ironic utterances and implicatures in terms of how much interpretive effort they require and of how radically they depart from encoded content. Given that metaphor (like conventionality itself) ranges across such a seamless continuum, we want to know what metaphors in general have in common. A bifurcated theory is in principle possible, on which only the most highly routinized but still not conventionalized metaphors count as part of 'what is said'. But this will leave out most of what we normally think of as metaphor. Neither Bezuidenhout nor Recanati suggests such a bifurcated account, and it seems clear that we should prefer a unified theory if possible.

4. 'What is Said' as the Vehicle for Implicature

The third criterion offered to mark the boundary of 'what is said' aims, like the second, to establish that 'what is said' comes first in the order of interpretation. However, this criterion invokes intuitions about which interpretations serve as inputs for which others, rather than intuitions about 'directness' *per se*. Bezuidenhout (2001b, p. 161; following Tsohatzidis, 1994) offers the following pair of examples to show that metaphorical interpretations can serve as 'springboards' for further interpretation:

- (19) Our piglet is getting dirty.
- (20) She's the Taj Mahal.

Bezuidenhout imagines (19) being uttered by a husband to bring his wife's attention to their child, in order to get her to stop the child's splashing around in the mud; she imagines (20) being uttered sarcastically in response to the question of whether the woman being discussed is attractive.

The fact that an interpretation can serve as a 'springboard' for further interpretation can't by itself mark the boundary of 'what is said', because both sarcasm and implicature can do this as well. For instance, suppose that Bill asks Alice whom they should invite for dinner; Alice could respond with a sarcastic utterance of:

(21) Well, Jane is always so utterly charming,

in order to implicate that Jane should absolutely not be invited. Similarly, suppose that Bill and Alice have a long-established routine of going to the movies on Thursday nights, and that the following exchange takes place between them early one week:

(22) Alice: So, which movie we should go see on Thursday?

Bill: Actually, I was thinking I might grab a beer with the guys from work.

Bill's utterance responds to Alice's question by implicating that he doesn't plan to go to the movies on Thursday (and so has no reason to offer an opinion about which movie they should see). At the same time though, the obvious lameness of Bill's excuse combined with mutual knowledge of their routine, further implicates that he is intentionally breaking that routine, and perhaps even breaking up with Alice altogether.

When can an indirect interpretation set up a further implicature? The operative requirement seems to be the same as we found in §2: in order for an implicature *Q* to be launched from an interpretation *P* of an utterance *U*, the speaker's intention for *U* to be interpreted as *P* has to be open and obvious, and not merely insinuated. Metaphor and sarcasm meet this requirement because they are such well-established routes for communicating something by an utterance other than its conventional meaning. (In the case of sarcasm, the distinctive tone of voice typically provides an explicit indication of the speaker's interpretive intention.) But an implicature can be launched from *P* whenever the speaker's intention for *U* to be interpreted as *P* is sufficiently obvious, even when *P* is itself an implicature. When this condition of obviousness fails—as with highly poetic, allusive metaphors like (11)—then *U* cannot launch a further interpretation.

However, Bezuidenhout also points out that when metaphor and sarcasm are combined, as in the imagined utterance of (20), then the sarcastic interpretation seems invariably to be *conditioned upon* the metaphorical one. If we give a rational reconstruction of the connection between the words uttered and the content ultimately communicated, we always appeal to the metaphorical interpretation

first. This remains true even when metaphor and sarcasm are combined with implicature—if, for instance, the speaker were to implicate by uttering (20) that he was unwilling to go out with the woman being discussed. Bezuidenhout concludes:

Thus, we can say that metaphors ... must be launched from an utterance and cannot be launched from an interpretation that is itself reached only via the pragmatic interpretation of a speaker's utterance Metaphorical interpretations have a kind of directness that is not shared by other pragmatically derived interpretations, such as cases of irony and indirect speech acts (2001b, pp. 163–4).

Why should the order of interpretation be fixed in this way? I find a suggestion of Josef Stern's (2000, p. 237) appealing: metaphor operates on expressions to determine propositional contents, while irony operates on propositional contents to determine new contents. (In the case of sarcasm, which is one species of irony, it selects the salient proposition from a set of 'contraries' to the input proposition.) If we grant that the interpretation of metaphor operates on expressions and that of irony on contents, then we can see why the order of interpretation should be fixed as it is: all word-based interpretation should ideally be completed before content-based interpretation begins. Beginning with the sarcastic interpretation in an utterance which combines metaphor and sarcasm would require interpreting the sentence uttered literally, then interpreting the resulting proposition sarcastically, and finally going back to the words uttered to interpret *them* metaphorically. This would make no sense: in particular, instead of one complex interpretation, we would be left with two distinct, parallel interpretations.

The crucial question, though, is whether granting this distinction between the two modes of interpretation shows that metaphor belongs within 'what is said', as Bezuidenhout and Stern maintain. I don't think it does. For there are precisely analogous complex utterances involving manner implicatures. Consider, for instance, a likely utterance of:

- (23) George's elocutionary style has been subjected to the most rigorous refinement by sustained exposure to the very best company.

Here, 'the very best' is intended to be interpreted according to the standard invoked by the utterance's elevated manner, and in particular by the use of 'elocutionary style' rather than a more prosaic, truth-conditionally equivalent expression, like 'way of talking'. By employing this mode of expression, the speaker implicates—but merely implicates—that the relevant sort of company is that of elegant, refined, upper-class people, and that the relevant sort of 'rigorous refinement' was learning to dish out erudite *bon mots* at endless dinner parties. If this utterance is intended sarcastically, then all of these implications will be included within the scope of sarcasm: the speaker may implicate that 'the very best company'

is lame and stuffy, and ‘rigorous refinement’ a boring waste of time.⁹ And given such a sarcastic interpretation, the speaker’s utterance might well implicate in turn that George, being so fancy, wouldn’t fit in with genuinely good company—say, with the guys who play pool down at the bar.

Contrast this with the case where the speaker utters (20) metaphorically and sarcastically, to implicate that he wouldn’t be willing to go out with the woman under discussion. In both cases, an interpretation *P* is first generated from the particular expressions employed in a sentence *S*; *P* is then interpreted ironically, producing interpretation *Q*; and this in turn generates a further implicature *R*. Both cases also exhibit the same constraint on the order of interpretation. (23) can’t be heard sarcastically as asserting the contrary merely of the propositional content of what was actually said, while also retaining a straightforward commitment to the implicatures generated by the elevated manner of speaking; the manner-generated implicatures must fall within the scope of the sarcasm too. As with metaphor, this seems to result from the fact that sarcasm operates on contents and so can swing into play only after all interpretations that depend on the particular expressions employed have been calculated. But again, we wouldn’t want to conclude from this that all interpretations which depend on the particular expressions employed belong on the ‘saying’ side of the divide.

5. ‘What is Said’ as What is Available for Explicit Response

The fourth and final argument for treating metaphor as part of ‘what is said’ is that metaphorical meaning is explicit, in the sense that later speakers can respond to metaphorical contents by echoing the original speaker’s own words. For instance, Bezuidenhout (2001b, p. 157) says about her ‘bulldozer’ dialogue (example (5) above) that it is natural to describe the later speakers as ‘either agreeing or disagreeing with what B says.’ David Hills (1997) makes much the same argument more explicitly. He points out that if someone responds to Romeo’s utterance of:

(24) Juliet is the sun

by saying either ‘No, she isn’t’ or ‘She sure is,’ then this is naturally construed as a response to Romeo’s metaphorical claim: that Juliet is beautiful, nurturing, worthy of worship, and so on. By contrast, the same sort of response to an utterance of (4) will accept or challenge what the speaker *said*—that John’s command of English

⁹ On the other hand, if George were obviously and saliently a coarse, loud, foul-mouthed sailor, then the sarcasm might implicate that George speaks in a decidedly unelevated tone. But here again the manner-generated implicatures will be included within the sarcasm, though in a rather different way: George has most assuredly not been hanging out at fancy dinner parties and dishing out *bon mots*.

is excellent and his attendance at tutorials regular—rather than the implicature that John is a bad philosopher. Hills concludes:

So it would appear that Romeo's meaning gets lodged in Romeo's words in a way that Grice's meaning (in the letter of recommendation example) never gets lodged in Grice's words. The words of Romeo's utterance, as used by him on a particular occasion, get taken so as to express a thought they wouldn't express if they were taken literally—one which may be true or false or indeterminate in its truth value, one to which we are free to respond in ways that are appropriate only to thoughts that speakers have actually put into words (1997, p. 127).

In short, because later uses of words like 'bulldozer' and 'sun' inherit the original speaker's metaphorical meaning, and so can be used to respond to that speaker's claim, it seems that we have explicit, systematic usage of a sort that justifies treating those words at least as playing the role of expressing *ad hoc* concepts, and perhaps even as having taken on new meanings, within that conversational context.¹⁰

Though this criterion is more promising than the others, I think it too fails to delineate an acceptable boundary for 'what is said', because it also applies to implicatures and other forms of speaker meaning, including sarcasm. Consider

- (25) Alice: John has three children.
Bill: No, he doesn't; he has four.
- (26) Alice: Some students came to class today.
Bill: No, they were all there.
- (27) Alice: John is happy.
Bill: No he's not; he's ecstatic.

In each case, Bill's utterance denies merely pragmatically conveyed content: that John has *at most* three children; that some *but not all* of the students were present; that John is *merely* happy. The added content must be merely pragmatic, because Alice could consistently have added Bill's second clause to her own, as in 'John is happy; in fact, he's ecstatic' (cf. Levinson 2000, 210–3). So the availability of this

¹⁰ Hills goes considerably further than Bezuidenhout in this regard. Where Bezuidenhout classifies metaphor as a pragmatic phenomenon, albeit as one belonging within 'what is said', Hills advocates a semantic view:

In addition to metaphorical truth values belonging to sentences, there are metaphorical references belonging to subsentential expressions. If we think of metaphorical truth values as determined by metaphorical sentence contents (metaphorical thoughts), we can think of metaphorical references as determined by metaphorical expression contents (metaphorical senses). In this fashion, metaphorical contents become full participants in the familiar recursive rignarole of compositional semantics (1997, pp. 146–7).

form of response can't itself establish that metaphor is a semantic phenomenon, as Hills claims.

Examples like (25) and (26) won't yet convince contextualists like Bezuidenhout or Recanati, however: these are precisely the sorts of cases which motivate their view that distinctively pragmatic material 'intrudes' into 'what is said'. This position is considerably less plausible, I think, as applied to (27): while (25) and (26) do seem to involve implicit qualification much like that found in enrichment and loose talk, it's not nearly as intuitive that Alice implicitly meant by (27) that John was *merely* happy.¹¹

Further, sarcastic speech also allows hearers to take up the speaker's claimed content by echoing her words. Suppose Alice sarcastically utters something like:

(28) Jane's really been a fine friend to me in these last few weeks.

Then Bill can agree with Alice's intended claim, that Jane has been a lousy friend, by echoing Alice's words sarcastically himself:

(28') Oh yes, she sure has: just the sort of ally and boon companion that we all dream of.

Similarly, if Alice utters:

(29) All the brilliant theorists must have gone to lunch,

intending the noun phrase sarcastically, Bill can respond by picking up on her intended use, as in:

(29') No they haven't; they're all just too lazy to come to work in the morning.

Finally, the relevant pattern of response is even available with 'ignorant' speaker meaning. For instance, if Alice utters:

(30) Smith is making quite a mess of those leaves,

while Jones is saliently though erratically raking his lawn, she may mean to claim (among other things) that *that salient guy* is making a mess. And Bill may respond to her utterance with something like:

¹¹ A more plausible analysis of (27) is as metalinguistic negation; see Horn, 1985. This analysis is itself controversial, however; see Kempson, 1986, Levinson, 2000.

- (30') No, that's wrong. He has got a method: he's making a bunch of small piles.

Intuitively, though, what Alice has *said* concerns Smith, not Jones. We can also imagine similar interchanges involving malapropisms, as in:

- (31) Alice: I heard that after Jane insulted his brother, he began spewing vile epitaphs at her.
Bill: No he didn't; he just said that she was mean, which she is.

In all these cases, the original speaker employs a way of speaking which later speakers can continue to employ if they so choose. In some cases, like metaphor and sarcasm, both the original speaker and her respondents are fully aware that they are using words in a non-standard fashion. With malapropisms, the speaker is not aware of this. Later speakers may therefore feel that they would make themselves 'complicit' in linguistic error by taking up the original speaker's way of talking, and so they may preface their response with a repair—for instance, they may respond to (28) with something like:

- (30'') That's Jones, not Smith. In any case, he does have a method...

Even here, though, later speakers *can* accommodate or 'play along' with the original way of speaking in their own responses. Thus, if contextualists want to retain their original criterion, on which content that later speakers can respond to belongs within 'what is said', then they will need to include all these cases within 'what is said'. And if they do, it appears that they will lose the theoretically valuable distinction between what the speaker actually said and what she tried but failed to say.

It is equally if not more important, I think, that in none of these cases can later speakers be *compelled* to echo the speaker's own way of using her words: they are always entitled to insist upon a literal construal. For instance, Benvolio can reject Romeo's utterance of (24) by saying something like:

- (32) That's absurd: Juliet isn't the sun. It's up there, and she's a normal girl, here on Earth. It's older than the Earth, and she's barely 14 years old. And it's bigger than the Earth, while her waist is just 26 inches around.

According to our everyday patterns of conversation, Romeo can't just dismiss a literalist challenge like (32) by saying 'You're wrong: Juliet *is* the sun', or even 'That's not the context of interpretation I was imagining'. He must acknowledge the consequences that Benvolio mentions *as* consequences of (24), on pain of incompetence with the word 'sun'. And because he isn't willing to accept these consequences as being true of Juliet—after all, he doesn't actually *believe* that Juliet

is the sun or up in the sky—we expect him to admit that he didn't really mean what he originally said.¹² The same thing goes for loose talk as well (along with 'ignorant' speaker meaning). For instance, if Alice utters:

(33) France is hexagonal,

then Bill is entitled to respond with something like:

(33') That's false: its borders have lots of twists and turns.

In doing so, Bill points out a fact about France which entitles him to reject Alice's application of 'hexagonal' to it. Given Bill's insistence on precision, Alice's only option is to retract her original utterance, retreating to the weaker claim that France is roughly hexagonal.

The important point is not merely that later speakers in all these cases are free to switch to a literal use of the original speaker's words. Rather, the crucial point is this: if the original speaker's utterance had genuinely 'lodged' a new metaphorical meaning in the words uttered, or even just had established a new, temporary use for them, then that meaning should necessarily be inherited by any later use of those same words in that same context which *responds* to the initial claim. If multiple meanings or uses are available in a single conversational context, then they should be distinct, much like the multiple meanings of 'bank' or of 'mass', or distinct tokenings of 'that' which ostend distinct objects. Any 'response' that exploits an alternative meaning should be a non-sequitur. But in all these cases, a response that insists upon a literal meaning still counts as a response, not a non-sequitur. The reverse does not hold: a switch from a literal to a non-literal use *is* a non-sequitur.

Of course, competent, cooperative, charitable hearers don't usually respond along the lines of (32) or (33'). They are usually much more interested in what the speaker does mean than in what she doesn't, and they naturally 'play along' by responding to that. This is especially palpable for utterances like (24): precisely because Romeo's utterance would be so absurd if construed literally, hearers automatically search for an alternative interpretation. Insisting upon the literal interpretation of an utterance like (24) is indeed non-cooperative in some sense of the term. But this sort of non-cooperativeness cannot impugn the challenger's competence as a linguistic interpreter—it is much more like the non-cooperativeness of refusing to indulge an eager child in a game of make-believe, or of refusing to

¹² A language may lack the expressive resources necessary for the speaker to state her intended meaning literally and explicitly (Camp, forthcoming). In such a case, the speaker may retreat to a simile, saying something like 'Fair enough. I meant that Juliet is *like* the sun with respect to the role she plays in my life; but I can't express exactly what that role is'. The fact that there may be no possible paraphrase of what the speaker did mean does not show that the speaker must have meant what she actually did say.

go along with someone's talk of ghosts or auras. Autistic and schizophrenic people, as well as small children, tend to insist on literalist interpretations of metaphorical utterances (cf. e.g. Langdon *et al.*, 2002). These people are certainly impaired *qua* conversationalists, but we don't usually think of them as suffering from a distinctively linguistic impairment, and we do normally feel compelled to acknowledge their interpretations as legitimate.

The legitimacy of a challenging literal response to a metaphorical utterance is clearer when the literal meaning is less absurd. Imagine that Socrates has been heard praising Diotima's virtues in the marketplace, uttering things like:

(35) Diotima is a midwife.

Suppose also that the city council is worried that new babies aren't being officially registered, and so they're threatening midwives with imprisonment if they don't register with the council. Hearing Socrates utter (35), a councilman could legitimately take him to have said that Diotima is a midwife, and could demand to know why Diotima hasn't registered. But suppose that Socrates meant (35) metaphorically: he meant that Diotima helps young men to work out their philosophical views on the Good. Then it would be overwhelmingly natural for him to respond to the councilman by saying something like:

(36) I didn't mean that she's really a midwife—I was just speaking metaphorically.

Even if he doesn't say this, Socrates can't just insist on his preferred interpretation while simultaneously denying that Diotima delivers babies. If he continued to utter (35) without a concessionary remark like (36), then the council would be entitled either to use his utterance as evidence against Diotima, or else to treat Socrates as incompetent with the term 'midwife', unless he retracted his utterance.¹³

A literal speaker cannot be forced to retract her utterance in this way. She can accept all the obvious consequences a hearer brings up, or else deny that they follow from what she said, without sacrificing her claim to linguistic competence. This is true even for utterances containing ambiguous or context-sensitive terms. For instance, suppose Bill is accused of robbing a bank, and Alice is reported as having uttered:

(38) Bill was planning to head over to the bank around noon.

¹³ If a speaker prefaces her utterance with something like 'Metaphorically speaking', then her hearers cannot insist upon a literal challenge. But they can insist that they don't understand what the speaker could mean, without impugning their interpretive abilities as they would if they insisted that they didn't understand a literal utterance. (Indeed, this posture of insistent 'metaphorical deafness' seems to be somewhat common among philosophers.)

Asked by the detective whether she said that, Alice would need to admit that she did, *and* that she meant it, but she could also add that she said that Bill was going to the riverbank, not to the financial institution. In this case, Alice could continue to utter (38) even after the detective's challenge, and no one could hold her responsible for placing Bill at the scene of the crime, or treat her as less than fully competent linguistically. Similarly, suppose that the university is trying to improve its sports programs, and Alice is heard to say:

(39) We should recruit Charles—he's tall.

On this basis, Bill the basketball coach proposes a lavish scholarship for Charles. Later in the conversation, it is discovered that Charles is just 5'10". Bill angrily accuses Alice of having made a mistake, because she said that Charles was tall. But if we suppose that Alice had originally been talking about Charles' potential as a gymnast, then she is perfectly entitled to stand by her original claim. *She* didn't make any mistake, because Charles is tall for a gymnast. It was Bill who misinterpreted her by taking the relevant standard to be that of basketball players.¹⁴

In the case of semantic context-sensitivity and ambiguity, then, the speaker is both prepared and entitled to stand by her original utterance, and the appropriate response to a challenge is to clarify how the challenge is a non-sequitur given one's intended reading. In metaphor and loose talk, by contrast, the speaker is not inherently entitled to her original utterance, and we expect her to respond to a challenge by retracting her utterance and replacing it with a more explicit statement of what she did mean. In this sense, speakers who don't say what they mean rely upon their hearers' indulgence *not* to hold them responsible for the literal meanings of their words. (Of course, all things considered, a non-literal way of speaking may be the best, most appropriate means for expressing one's intended meaning. All things considered, it may be best to lie, too. Conversation involves multiple, possibly conflicting norms.) Even in contexts where hearers do accommodate the speaker's intended meaning, the literal meaning retains a sort of normative priority over the non-literal use.

¹⁴ Cf. Saul, 2002, pp. 351-2. There are extreme cases in which contextual factors so strongly support the hearer's interpretation that he could not reasonably be expected to recover the speaker's intended meaning. Suppose the conversation to this point has revolved entirely around financial institutions, and the speaker says 'I'll meet you at bank', meaning the riverside. The speaker can insist that he meant 'bank' in the riverside sense, but the hearer could legitimately object that there's no way he could have known this. The unrecoverability of the speaker's intended meaning must be quite dramatic, though, for such an objection to be appropriate. Indeed, it's plausible that a reasonable speaker could not even intend a reading is so obviously unrecoverable.

6. What is 'What is Said'?

Contextualists have assumed that we could redraw the boundary of 'what is said' to accord more fully with ordinary intuition while also retaining a traditional notion of implicature. Indeed, they thought we needed to redraw that boundary precisely in order to respect the traditional notion: so as not to conflate implicature with pervasive, more or less unreflective phenomena like loose talk. However, all of the criteria they offer to delineate the boundary of 'what is said' have either supported the traditional boundary or else have marked a distinction among what even contextualists agree to be implicatures.

Terminological issues aside, the important point is that our ordinary practices of conversational exchange require some notion of 'first meaning'—first in the rational order of interpretation, not necessarily in the process of actual comprehension—that captures the content of the speaker's locutionary act. These same practices justify excluding metaphor from this 'first meaning'. Ordinary speakers do recognize the standards for literal use that are built into words' conventional meanings. They do think that metaphorical meaning is indirect, in the sense of exploiting these literal meanings to an alternative end. And they do expect speakers to be able to justify, at least in rough outline, how their intended metaphorical meaning can be derived from the literal.

However, none of this yet tells us what *is* said by a given utterance. I agree with contextualists that our theoretical notion of 'what is said' should be grounded in pragmatic considerations about what speakers do in making their utterances. So I reject the traditional view, which stipulates that 'what is said' just is conventionally encoded semantic meaning. But I also believe that our notion of 'what is said' must be refined enough to provide a useful constraint on semantic theorizing about the conventional meanings of words and sentences. So I reject the contextualist view (which is also the view of skeptics about 'what is said' articulated by Cappelen and Lepore 1997) that 'what is said' is radically disconnected from conventional meaning. We lack pure, unmediated access to semantic meaning *per se*, and so we need some criterion for deciding which uses of language should serve as the data for semantic theorizing. Although this is obviously an enormous undertaking of its own, I want to conclude by sketching how, on the basis of the discussion to this point, we might begin to construct a theoretically workable notion of 'what is said' that meets these desiderata (see Figure 3). I don't want to claim that this notion captures *the* true meaning of 'what is said', let alone that we can or should ultimately reduce semantic meaning to the notion I sketch. Rather, I simply want to provide us with a more explicit and robust grasp on an intuitive and theoretically useful notion of 'first meaning' that I believe is already implicitly operative in our ordinary conversational practice.

In §2, I pointed out that it is normally only appropriate to report speakers as having 'said' contents to which they have openly and obviously committed themselves by their utterances. Because context-specific information and general principles of charitable interpretation contribute heavily to what it's reasonable to




Figure 3 *My View.*

treat a speaker as having obviously committed herself to, reports of ‘what is said’ in this sense typically depart significantly from semantically encoded content. However, this is not the only ordinary use of ‘say’. In §5, I relied on another use: the practice of holding speakers accountable for the consequences of what they said, strictly speaking.¹⁵ Of course, hearers are typically quite willing to accommodate non-literal talk; it is often by far the most efficient way to communicate one’s intended meaning. But speakers and hearers are also aware that such accommodation is merely optional. Speakers may feel that a hearer who insists upon a strict interpretation is being pedantic or silly, or is refusing to focus on the point at hand, but they also feel compelled to admit that, like it or not, that *is* a consequence of what they said.¹⁶

Both of these uses of ‘say’ share a focus on the commitments the speaker undertakes by virtue of making her utterance;¹⁷ I think we can bring them together to develop a useful notion of ‘what is said’. Begin with the fact that, within any given context, speakers and hearers share, or are at least able to coordinate upon, a relatively robust set of intuitions about what does and doesn’t obviously follow from what the speaker said. Sometimes things go wrong: a given speaker or hearer

¹⁵ Notice that this strict use is not simply quotational: for instance, Benvolio holds Romeo to substantive, empirical consequences of his utterance, literally construed. This strict but substantive use of ‘say’ is also operative in cancellation: a speaker disavows unwanted commitments that a hearer might reasonably and charitably derive from her utterance, but that are not actually required by it. However, speakers can also ‘cancel’ unwanted interpretations of ambiguous terms (cf. Sadock, 1978): for instance, a speaker can consistently say ‘Bill went to the bank, but I don’t mean to suggest that he went to the financial institution’. By contrast, focusing on what hearers are entitled to hold speakers responsible for does distinguish ambiguity from what is merely meant.

¹⁶ Lewis (1979, p. 352) claims that ‘the rule of accommodation [for standards of precision] can go both ways’. He immediately goes on to note, though, that ‘for some reason raising of standards goes more smoothly than lowering Raising of standards . . . manages to seem commendable even when we know that it interferes with our conversational purpose’. I suggest that the asymmetry arises because, while we often let each other ‘get away with’ low standards, there’s no need to ‘get away with’ a high-standards interpretation: a high-standards interpretation reflects the conditions of satisfaction specified by what the speaker actually said.

¹⁷ Cf. also Brandom, 1983; Green, 2000.

may have quite firmly held but idiosyncratic assumptions about what entails what. In general, though, within specific contexts we are exceptionally good at coordinating expectations about such consequences. This provides us with a very rough first approximation of 'what is said' by an utterance, articulated in (S₁). (I add the subscript 'τ' to make clear that this is a technical use of the phrase; we could also employ a different term entirely.)

- (S₁) *The content of what a speaker says_τ in uttering a sentence S in context C is fixed by the set of commitments O a competent hearer is entitled to hold the speaker accountable for having obviously undertaken by virtue of assertively uttering S in C.*

(S₁) appeals to assertive utterances of *S* in order to isolate 'what is said_τ' by an utterance of *S*. Like Austin (1962), Bach (1999), and other theorists, I think that saying is a locutionary act, while assertion is illocutionary: I think that speakers can assert things they do not say, and say things they do not assert. For these same reasons, one can assertively utter *S* without actually asserting *S*. However, I also think we can most effectively elucidate a useful theoretical notion of saying by focusing on the illocutionary act to which saying typically serves as a means. (I'm setting aside for now cases in which *S* is uttered without any illocutionary commitment, such as giving an example or reciting a play. I'm also focusing exclusively on assertion. Both restrictions would need to be removed in a full analysis.)

As stated, (S₁) includes *all* the commitments the hearer is entitled to hold the speaker accountable for. We can start to winnow down the commitments in *O* by examining it for consistency. When a speaker says something that she doesn't mean in order to mean something else, *O* will include conflicting consequences:¹⁸ some of the commitments derived from a charitable interpretation will conflict with those derived from a strict interpretation. By contrast, there is typically no such conflict when a speaker means what she says, even if she also means something else. The fact that *O* is inconsistent thus serves as strong evidence that the speaker can't have meant what she actually said. This evidence is not decisive, of course: a speaker may really mean something that conflicts with the obvious facts, or that is logically inconsistent. To the extent that there is an obvious conflict between obvious consequences, though, the hypothesis that the speaker does not mean what she said is rendered more plausible. We can rule out most metaphor and sarcasm from 'what is said_τ' on this basis.¹⁹

¹⁸ Assuming the utterance is not so defective as to render the additional meaning unrecoverable.

¹⁹ No conflict is generated for 'twice true' and 'twice apt' metaphors. With a 'twice true' metaphor, the speaker believes the commitments generated by the strict interpretation; she merely doesn't intend them to be registered as conversational consequences of her utterance. In the case of 'twice apt' metaphors, this latter condition is met as well.

We can also exploit the conflict to narrow down the commitments in *O*. Some of the commitments in *O* are merely optional: they follow from charitable assumptions about why a speaker might make this sort of utterance. Other commitments are obligatory: upon being challenged, the speaker must admit that those are consequences of assertively uttering *S*, like it or not. This fact, that some consequences are merely optional while others are obligatory, itself establishes a hierarchical distinction between the two classes of commitments. When there is a conflict between commitments, where some are optional and others obligatory, only the obligatory ones should be retained as elements of the refined version of *O*. (If they are both obligatory, retain both).²⁰

Of course, it may be obvious from independent facts about the conversation that, all things considered, a hearer should not actually attribute all of the commitments in the refined set to the speaker. Even then, though, as I argued in §5, the hearer is still entitled to hold the speaker accountable for them.²¹ Even when a hearer has good overall interpretive reasons not to hold the speaker accountable for certain commitments in the obligatory class, he cannot be compelled to disregard those commitments. Nor can he be obliged to accept merely optional commitments. He can insist on flat-footedly treating the speaker as committed to the obligatory consequences unless and until the speaker explicitly disavows her utterance. By contrast, the speaker is obliged to acknowledge the obligatory commitments; she can only dislodge them by disavowing her original utterance.

These considerations suggest a revised version of (S₁):

- (S₂) *The content of what a speaker says_T in uttering S in context C is fixed by the set of commitments O which a competent hearer is entitled to hold the speaker accountable for having obviously undertaken by virtue of assertively uttering S, and which the speaker is obliged to acknowledge as consequences of having assertively uttered S, in C.*

Though much more restrictive, (S₂) is still too encompassing. Consider a context in which a speaker says *S* and means it: she willingly accepts all the obvious, obligatory consequences a hearer can legitimately attribute to her. These consequences will almost certainly still differ across different contexts of utterance, because the fact that the speaker assertively uttered *S* will interact with conversational presuppositions to produce obvious, obligatory commitments that would not follow from an utterance of *S* in a conversation involving different presuppositions.

²⁰ In effect then, I am suggesting that we can develop a theoretical notion of 'what is said_T' by exploiting a pattern of conversational behavior which is crucially associated with an ordinary strict use of 'say', but which is in principle distinct from brute intuitions about the extension of 'what was strictly speaking said'

²¹ Likewise, some of the commitments included in *O* may already have been undertaken by the speaker in virtue of earlier conversational moves. But a hearer is still entitled to cite the speaker's utterance of *S* as evidence that the speaker has undertaken these commitments.

Recall the example in which a professor responds to the question of whether Alice passed her exam by saying:

(8) I didn't fail any students.

So long as it is an obvious shared conversational assumption that Alice took the exam, and that every student either passed or failed—we can even suppose that those assumptions have been articulated explicitly earlier in the conversation—the claim that Alice passed will be included within the set of commitments O generated by (S_2) . In such a context, the professor would normally be obligated to accept this commitment as a consequence of asserting (8).

To get a workable definition of 'what is said,' then, we still need to weed out the effects produced by the additional background presuppositions of particular conversations. The natural way to do this is to consider various conversational contexts in which a given sentence is uttered, and to include only obligatory commitments that would be obviously undertaken by the same utterance-type across contexts. We need to include both actual and possible contexts of utterance to account for 'what is said_T' by utterances of sentences which can't in fact be uttered, for instance because they are too long. However, we cannot include absolutely all possible contexts of utterances, because bizarre contexts—such as those in which speakers and hearers happen to share wildly idiosyncratic assumptions about word meaning—will swamp the analysis. At the same time, if 'what is said_T' is to help us isolate sentences' literal, semantic meanings, we cannot appeal directly to just those contexts in which words are used literally and with their conventional meanings.²² Therefore, I suggest that we should quantify over 'normal contexts of utterance', where normalcy is cashed out in sociological rather than semantic terms. Normal contexts of utterance, very roughly, are contexts involving conversational exchanges between people who are treated by their community as competent speakers of English, and where no explicit, special suppositions about word meaning have been established. This gives us (S_3) :

(S_3) *The content of what a speaker says_T in uttering S in context C is fixed by the set of commitments O a competent hearer could hold the speaker accountable for having obviously undertaken by virtue of assertively uttering S, and which the speaker would be obliged to acknowledge as consequences of having assertively uttered S, in any normal context in which a speaker assertively uttered S.*

²² Scott Soames makes a related proposal in chapter 3 of Soames, 2002, although he allows himself to appeal explicitly to normal contexts in which speakers use words with their literal meaning and do not speak metaphorically, sarcastically, etc. Soames is concerned to defend a particular, Millian view of names, which raises problems for his analysis of attitude ascriptions. By contrast, I have no antecedent commitments to the proper analysis of any particular class of terms (though of course (S_2) itself does rule out some analyses and rule in others). Thanks to Mark Richard for pressing the point about attitude ascriptions.

This reduces *O* dramatically. However, (*S*₃) still includes presuppositions which would be triggered by the specific sentence as uttered in any normal context of utterance. We can isolate *these* commitments by testing each member of *O* to see whether it would be included in the analogous set of commitments generated by an assertive utterance of the negation of *S*. Call the set of commitments which pass this test *P*. While the members of *P* presumably do need to be subtracted from *O*, we cannot simply disregard them in the way we could disregard context-specific presuppositions included in (*S*₂). Members of *P* are relevant to semantic theorizing in a way that contextual presuppositions are not.²³

If we adopt (*S*₃), then determining what is said_T by any given utterance will require investigating utterances of that same sentence in other contexts. This is indeed a theoretical project: as Cappelen and Lepore (1997) and Scott Soames (2002, p. 68) say, ordinary speakers don't have intuitions about the commitments invariantly undertaken across contexts. But ordinary speakers do have intuitions, however unrefined, about commitments undertaken *in* contexts. The theoretical challenge is to generalize appropriately across these intuitions. Though this project is not easy, I believe that it is feasible, and that it is sufficiently grounded in ordinary intuitions about what speakers actually do in making their utterances to have real empirical traction. I also think that it is a project with sufficiently close ties to semantic theory. A semantic theory makes predictions about a sentence's meaning on the basis of what it postulates to be the conventional meanings of its constituent terms and their mode of combination. We normally test such theories by employing something very like (*S*₃): by trying to imagine plausible scenarios in which speakers assertively utter a given sentence but should not be held accountable for the predicted content. The absence of such counter-examples supports the theory's claim to capture 'what is said_T' by an assertive utterance across contexts.

Obviously, though, (*S*₃) falls far short of an adequate theory of 'what is said.' Let me flag just a few of the most glaring issues. First, (*S*₃) says that the content of what is said_T in uttering *S* in *C* is 'fixed by' commitments obviously undertaken across

²³ Jason Stanley has objected (in conversation) that testing for presupposition will block a correct analysis of 'what is said_T' by some sentences. Consider the sentence 'An utterance has been made': it looks like the set *P* will include a commitment to the proposition that an utterance has been made, on the ground that it's undeniable that an utterance *has* been made even when the uttered sentence is 'It is not the case that an utterance has been made'. However, I believe that ordinary speakers do intuitively distinguish between commitments undertaken by virtue of *what* the speaker said, and commitments undertaken by virtue of *the fact that* the speaker said it; and I think we can legitimately avail ourselves of this distinction to make sense of such examples. Indeed, it sounds quite odd to say that by uttering ' $2 + 2 = 4$ ', the speaker has *undertaken a commitment* to an utterance having been made, even though the speaker cannot very well deny that an utterance has been made. In particular, the speaker is under no obligation to defend this commitment or else retract his original utterance. (Cf. also García-Carpintero, 2001, p. 102.)

contexts. I have not specified how contents are fixed by commitments. Roughly, when we commit ourselves to *things being a certain way*, that 'way' is the content of our commitment. However, it is an open question which order of explanation and which criteria we should employ in specifying the relation between contents and commitments. Further, as Soames (2002, pp. 61–2) argues, we need to impose some explanatory order among the commitments generated, in order to focus in on the crucial commitments and to exclude trivial consequences from the content of 'what is said' itself. Finally, some of the commitments undertaken are not commitments to truth-conditional content, but concern appropriate affective attitudes, or the appropriate manner of speaking, or the current utterance's relation to previous ones. Some theorists think these commitments should be included within 'what is said', while others do not. I have said nothing to settle this matter—though I do take it to be an advantage of (S₃) that it enables us to include these factors within 'what is said' if we so choose.

Second, I have suggested that 'what is said' as defined by something like (S₃) should serve as a constraint on a theory of semantic meaning, and that we can thereby preserve a fairly clean, tight connection between the two notions. This leaves open many questions about just how tight that connection can and should be. To take one example, Grice's preferred understanding of 'what is said' distinguishes between active and passive forms of expression, because he includes 'the order in which [words] appear' as a factor in determining 'what is said.' On (S₃), the crucial question is whether active and passive forms of expression make the speaker liable for different commitments. They don't seem to. But a theory of semantic content might well predict that the two sentences' differing modes of combination produce different semantic contents. The more coarsely we individuate the set of commitments specified by (S₃), the more difficult it will be to retain compositionality, and the more slippage we may need to allow between it and semantic content. Some such slippage seems to me to be appropriate, but other cases may demonstrate that it is too great to be borne, and so that we should impose additional constraints on 'what is said_T'.

Third, and most crucially, I have not specified how to fix the contents of utterances containing expressions that are semantically context-sensitive. On (S₃) as it stands, there will be no single commitment, or only an extremely weak existential commitment, undertaken by all utterances of a sentence like:

(40) That's my favorite song these days,

given the wide variety of possible speakers, ostended songs, and times of utterance. This is obviously unsatisfactory. To arrive at the appropriate content for utterances containing context-sensitive terms, we need to modify (S₃) so that it extracts a general commitment-schema for utterances of a given sentence-type containing one or more semantically context-sensitive terms. This general schema will in turn determine token commitments for a given utterance of that sentence in a specific context. I have said nothing about how this should be done.

This returns us to the central question surrounding contextualism: when we should count a term as semantically context-sensitive, and when we should attribute the variation in contribution to communicated content as due either to ambiguity or else to pragmatic factors. Roughly, with ambiguity we find that the contents contributed by a given term across all contexts systematically cluster around a limited number of distinct options, with little or no systematic relation among those options. In the case of semantic context-sensitivity, we find greater variability in the contents contributed, but we can isolate a constant, common schema for determining particular contents in particular contexts. Finally, and crucially, in the case of mere pragmatic influence, speakers are not treated as necessarily entitled to their intended content. Such speakers are vulnerable to the sorts of challenges cited in §5, and are expected to be capable of providing the sorts of rational reconstructions of their intended meaning cited in §3.

We might well worry whether the differences among these patterns of challenge, accommodation, and justification are sufficiently robust to draw the boundary in what most traditional semanticists would consider to be the 'right' place: to show that, say, the intended contents of polysemous terms and graded adjectives do belong within 'what is said', while loose talk and meaning enrichment do not. The data are less decisive for loose talk than they are for metaphor and irony, but I think they still clearly support excluding loose talk from 'what is said'. I am not so sure about meaning enrichment: I find the contextualist arguments most convincing here. I suspect that (S₃) itself is not robust enough to decide these hard cases. Hopefully we can appeal to independent tests for semantic context-sensitivity, such as bindability (Stanley, 2000), to adjudicate some of these debates. Perhaps admitting a category of what Bach (1994) calls 'implicature'—cancelable contributions to utterances' truth-conditions—can bridge some of the gap between a comparatively minimalist notion of 'what is said' and what we intuitively think of as implicature. And perhaps we need to allow some pragmatic 'intrusion' into 'what is said' after all—albeit much less than many contextualists currently envision.

Although these are obviously large and pressing questions, answering them is a task for another time. My primary aim has been to argue that metaphorical meaning should not be included within 'what is said'. Contextualists are correct, I think, that our theoretical notion of 'what is said' should be grounded in ordinary intuition. And they are correct that saying involves the direct, explicit expression of one's intended meaning. But they are wrong that metaphor counts as saying in this sense. Metaphor does not meet the criteria offered by contextualists themselves for identifying 'what is said' by an utterance.

Because philosophers of language have recently paid so much attention to the effects of context on semantically encoded content, we are now much more aware of the many varieties of contextually-sensitive linguistic phenomena which fall between the two extremes of heavily rule-bound, conventionally context-sensitive terms like 'I', and 'insinuating implicatures' like Grice's letter of recommendation. This range includes (at least) 'parametric' semantic context-sensitivity, enrichment, loose talk, 'nonce sense' (Clark, 1983), standardized indirect speech, and figuration. We should

continue to pay close attention to the particular behaviors manifested by these different phenomena. We need an overall theoretical classification which respects the significant differences among them. Perhaps we will find reasons to include some of these phenomena within the realm of 'what is said,' and even of semantics. Metaphor, however, will almost certainly not be among this class, for the reasons cited above.

*Society of Fellows
Harvard University*

References

- Austin, J. L. 1962: *How To Do Things with Words*. Cambridge, MA: Harvard University Press.
- Bach, K. 2001: Speaking loosely: sentence nonliterality. *Midwest Studies in Philosophy*, 25, 249–263.
- Bach, K. 1999: The Semantics-pragmatics distinction: what it is and why it matters. In K. Turner (ed.), *The Semantics/Pragmatics Interface From Different Points of View*. New York: Elsevier.
- Bach, K. 1994: Conversational implicature. *Mind & Language*, 9, 124–62.
- Bezuidenhout, A. 2001a: Radical pragmatics. In M. O'Rourke, D. Schier, and J. Campbell (eds), *Topics in Contemporary Philosophy: Meaning and Truth*. New York: Seven Bridges Press.
- Bezuidenhout, A. 2001b: Metaphor and what is said: a defense of a direct expression view of metaphor. *Midwest Studies in Philosophy*, 25, 156–186.
- Blasko, D. and C. Connine 1993: Effects of familiarity and aptness on metaphor processing. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 19, 295–308.
- Bowdle, B. and D. Gentner 2005: The career of metaphor. *Psychological Review*, 112, 193–216.
- Brandom R. 1983: Asserting. *Noûs*, 17, 637–650.
- Brisard, F., S. Frisson and D. Sandra 2001: Processing unfamiliar metaphors in a self-paced reading task. *Metaphor and Symbol*, 16, 87–108.
- Camp, E. 2006: Metaphor and that certain 'je ne sais quoi'. *Philosophical Studies* (forthcoming).
- Cappelen, H. and E. Lepore 1997: On an alleged connection between indirect speech and the theory of meaning. *Mind & Language*, 12, 278–296.
- Carston, R. 2002: *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Carston, R. 1999: The semantics/pragmatics distinction: a view from relevance theory. In K. Turner (ed.), *The Semantics/Pragmatics Interface from Different Points of View*. London: Elsevier Science.
- Clark, H. 1983: Making sense of nonce sense. In G.B. Flores d'Arcais and R. J. Jarvella (eds), *The Process of Language Understanding*. London: Wiley.
- Cohen, T. 1976: Notes on Metaphor. *Journal of Aesthetics and Art Criticism*, 34, 249–259.
- Crimmins, M. 1998: Hesperus and Phosphorus: sense, pretense, and reference. *Philosophical Review*, 107, 1–47.
- Davidson, D. 1986: A nice derangement of epitaphs. In E. Lepore (ed.), *Truth and Interpretation: Perspectives on the Philosophy of Donald Davidson*. Oxford: Blackwell.

- García-Carpintero, M. 2001: Gricean rational reconstructions and the semantics/pragmatics distinction. *Synthese*, 128, 93–131.
- Giora, R. 2003: *On Our Mind: Salience, Context, and Figurative Language*. Oxford: Oxford University Press.
- Giora, R. 2002: Literal v. figurative language: different or equal? *Journal of Pragmatics*, 34, 487–506.
- Giora, R. 1997: Understanding figurative and literal language: the graded salience hypothesis. *Cognitive Linguistics*, 8: 183–206.
- Green, M. 2000: Illocutionary force and semantic content. *Linguistics and Philosophy*, 23, 435–473.
- Grice, H. P. 1975: Logic and conversation. In P. Cole and J. L. Morgan (eds), *Speech Acts: Syntax and Semantics vol. 3*. New York: Academic Press.
- Grice, H. P. 1989: Utterer's meaning and intentions. In *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Hills, D. 1997: Aptness and truth in verbal metaphor. *Philosophical Topics*, 25, 117–153.
- Horn, L. 1985: Metalinguistic negation and pragmatic ambiguity. *Language*, 61, 121–174.
- Kempson, R. 1986: Ambiguity and the semantics–pragmatics distinction. In C. Travis (ed.), *Meaning and Interpretation*. Oxford: Blackwell.
- Langdon, R., M. Davies and M. Coltheart: Understanding minds and understanding communicated meanings in schizophrenia. *Mind & Language*, 17, 68–104.
- Levinson, S. 2000: *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. Cambridge MA: MIT Press.
- Lewis, D. 1979: Scorekeeping in a language game. *Journal of Philosophical Logic*, 8, 339–359.
- Lewis, D. 1983: Languages and language. In *Philosophical Papers, Vol. 1*. Oxford: Oxford University Press.
- Martinich, A. 1984: A theory for metaphor. *Journal of Literary Semantics*, 13, 35–56.
- Perry, J. 1986: Thought without representation. *Proceedings of the Aristotelian Society*, Supp. Vol. 60, 263–283.
- Recanati, F. 2004: *Literal Meaning*. Cambridge: Cambridge University Press.
- Recanati, F. 2001: Literal/nonliteral. *Midwest Studies in Philosophy*, 25, 264–274.
- Recanati, F. 1995: The alleged priority of literal interpretation. *Cognitive Science*, 19, 207–232.
- Sadock, J. 1978: On testing for conversational implicature. In P. Cole (ed.), *Syntax and Semantics*, Vol. 9: Pragmatics. New York: Academic Press.
- Saul, J. 2002: What is said and psychological reality: Grice's project and relevance theorists' criticisms. *Linguistics and Philosophy*, 25, 347–372.
- Searle, J. 1979: Metaphor. In *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge: Cambridge University Press.
- Soames, S. 2002: *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford: Oxford University Press.
- Stanley, J. 2000: Context and logical form. *Linguistics and Philosophy*, 23, 391–434.
- Stern, J. 2000: *Metaphor in Context*. Cambridge, MA: MIT Press.
- Tsohatzidis S. 1994: Speaker meaning, sentence meaning and metaphor. In S. L. Tsohatzidis (ed.), *Foundations of Speech Act Theory: Philosophical and Linguistic Perspectives*. London: Routledge.