

THE MEANING OF WHAT IS SAID AND  
THE ASCRIPTION OF ATTITUDES

# THE SEMANTIC SIGNIFICANCE OF WHAT IS SAID

*Emma Borg*

## *Abstract*

*It is often held that a correct semantic theory should assign a semantic content,  $p$ , to a given sentence,  $s$ , just in case a speaker who utters  $s$  says that  $p$  – thus ‘what is said’ is taken to be a semantically significant notion. This paper explores what exactly such a claim amounts to and offers five versions of the relationship between a semantic theory and judgements of what is said. The first three of these versions embody the central claim of semantic significance; however, I argue that none of these versions are feasible. Thus, contrary to the initial proposal, I claim that ‘what is said’ is not a semantically significant notion. Assignments of semantic content do not turn on evaluations of what a speaker uttering a sentence says.*

The question of how to correctly analyse sentences of *oratio obliqua* (indirect reported speech) has received a good deal of attention within the confines of philosophy of language and one reason for this, we might think, is that such sentences may prove interesting not only in their own right, but also because of the role they play in informing us about the meaning of the sentences they report. If a speaker,  $U$ , who utters a sentence,  $s$ , can be correctly reported by an act of *oratio obliqua* of the form ‘ $U$  said that  $p$ ’ (where ‘ $p$ ’ supplies the content sentence of the indirect speech act), then it seems that we have learnt something about the meaning of  $s$ . Now, exactly how we spell out this thought is open to question (as we will see in §1), but the intuitive idea which it embodies is that there is a move to be made *from* reports of what is said *to* the contents of a correct semantic theory. Such a stance, which I label the ‘standard view’, has a distinguished history, but it is a position which I wish to reject in this paper. I will argue that there is *no* semantically privileged notion of ‘what is said’, and thus no considerations concerning *oratio obliqua* should constrain or otherwise affect our semantic theorising. In no case of indirect speech reporting is it possible to move from the acceptability of ‘ $U$  said that  $p$ ’ to a claim that  $p$  provides the correct semantic analysis of the original sentence  $s$ .

The structure of the paper is as follows: in §1 I will set out five distinct ways in which the relationship between semantic theories and indirect speech reports might be construed. The initial three positions, (a-c), embody the key

idea of semantically informative indirect speech acts, while the latter two, (d-e), do not. The simplest version of the standard view, given by (a), will quickly be seen to be implausible, but (b) and (c) initially remain live options. However, in §2, I will argue that this initial impression is mistaken; on closer inspection neither position is able to deliver the semantically informative notion of ‘what is said’ as promised.<sup>1</sup> Yet, I will suggest in §3, this finding would be problematic *only if* the sole alternative to the standard view were given by position (e), for it too is problematic. This is not the case, however, for the theorist who rejects the semantic relevance of ‘what is said’ may instead opt for the more moderate position, to be given by (d), whereby there is a one-way dependence *from* semantic analyses *to* indirect speech reports. This position both avoids the earlier objections and, I will argue, is independently attractive.

## 1. The Relationship Between a Semantic Theory and Indirect Speech Reports

Initially, at least, the idea that a semantic theory should tell us what information an utterance of a given sentence conveys seems quite unobjectionable, indeed it seems a mainstay of much formal semantic theorising. It seems entirely plausible to maintain that, if *U* can be reported as having said that *p* by her utterance of *s*, then *p* in some way gives the meaning of *s*. Although I don’t want to engage in an exegetical history of this position, I think it is not particularly controversial to label this stance the ‘standard view’.<sup>2</sup> However, recently such a view has come under attack, perhaps most notably in Cappelen and Lepore 1997, where it is objected that facts about reported speech in ordinary language do not support the claim that the notion has any significant role in influencing the construction of a semantic theory. They write: ‘That a semantic theory should specify what is said by utterances of sentences seems innocent enough, but, when this assumption is embodied

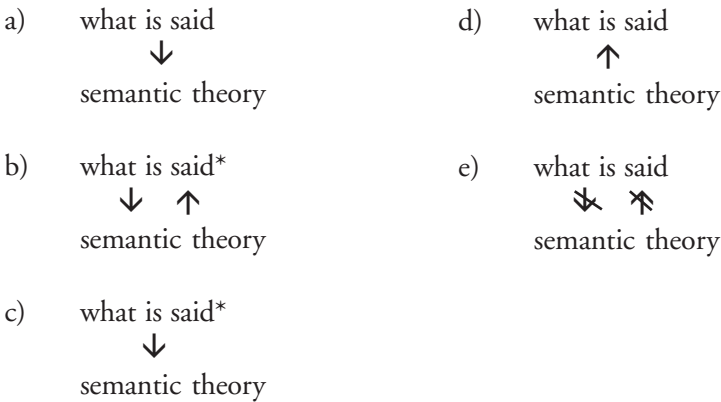
- 1 In what follows I will use the notions of reported speech and judgements of what is said interchangeably. This might be thought to be problematic since judgements of what is said range wider than acts of reported speech. For instance, an interlocutor may make a judgement of what is said by making a judgement of equivalence between two or more sentences, or by answering questions about cancelability, etc. (I’m grateful to Kent Bach for stressing this point). However, even in these latter cases, it seems that the agent puts herself in a position to make or endorse a report of the form ‘*U* said that *p*’, even if such an indirect speech report is not actually produced. Thus it seems to me that some equivocation between the two notions is not here pernicious.
- 2 The ancestral home of this view is probably the work of Paul Grice (cf. Grice 1989).

by MA [the principle that an adequate semantic theory for a language L should assign p as the semantic content of a sentence S in L iff in uttering S a speaker says that p], semanticists both misconstrue the aim of semantics and unreasonably constrain the semantics for indirect speech' (279). Ultimately, I want to side with Cappelen and Lepore as against what I have called the 'standard view', but to begin to get clear on what is at stake in the debate here, it seems that we need to start by establishing exactly what is being embraced by the former view and rejected by the latter – what is the relationship between a semantic theory and indirect speech reports that is under dispute?

I think there are (at least) five ways in which we might construe the relationship here (these are illustrated in Fig.1):

- a) One-way dependency - semantic analyses of a sentence, s, are dependent on the content of indirect speech reports concerning s. On this approach the content of a correct semantic clause for s is *constituted* (at least in part) by the content of correct indirect speech reports for an utterance of s.
- b) Two-way dependency: semantic analyses of a sentence, s, are dependent on the content of *appropriate* indirect speech reports ('what is said\*') concerning s, but which indirect speech reports are appropriate is itself a matter which is influenced by the semantic theory in play.
- c) One-way dependency: semantic analyses of a sentence, s, are dependent on the content of *appropriate* indirect speech reports ('what is said\*'), and which indirect speech reports are appropriate is determined by non-semantic features.
- d) One-way dependency - semantic analyses to indirect speech reports. Judgements concerning correct indirect speech reports are affected by the semantic analysis of the original sentence, but not vice-a-versa.
- e) No dependency: judgements of correct indirect speech reports and the output of a correct semantic theory are independent of one another – one concerns communication and the other linguistic meaning, and these are distinct domains.

Figure 1:



What unites (a-c), as against (d-e), is that there is held to be some kind of connection *from* (certain) indirect speech reports *to* a correct semantic theory. This supports the idea that indirect speech reports can be genuinely semantically informative or significant, for we can use the former to determine in some way the content of the latter. This is the claim made by the standard view; so is any such stance plausible?

Well, it seems fairly clear that the simplest construal of the standard view, given by (a), whereby fixing the facts about indirect speech reports fixes the facts for our semantic theory, is not plausible. For it is clear from even a momentary survey of the facts concerning reported speech in ordinary language that they diverge wildly from the kinds of features which can plausibly be taken to be the concern of semantics. To see this, imagine that Jim utters:

1) Blair lives at No.10.

(1) is a simple subject-predicate sentence, containing no overt indexicals, and thus would seem one of the most elementary cases for the advocate of the semantic significance of indirect speech reports to deal with. Yet even here the problem is immediately obvious, for the utterance of (1) does not stand in a single, one:one relationship with some indirect speech report, but rather explodes into a plethora of possible indirect speech acts, all of which may be licensed by Jim's production of (1). For instance, given the right sort of context of utterance and report, any of the following (amongst an indefinite number of others) may be acceptable:

- 2) Jim said that Tony Blair lived at No.10 Downing St, London, UK.
- 3) Jim said that the current Prime Minister lives at No.10.
- 4) Jim said that that man lives there.

- 5) Jim said that Baby Leo's father lives in No.10.
- 6) Jim said that the most right-wing Labour leader ever secured power where other, more left-wing, predecessors had failed.
- 7) Jim said that he knows about British politics.

Clearly, even in a basic case like (1), then, there are a vast range of potentially correct acts of reported speech, but intuitively we don't want to countenance *all* of them as semantically informative as to the meaning of the original sentence.

For one thing, allowing all indirect speech reports to be semantically relevant would mean rejecting the idea that a single proposition provides the content of a simple (non-indexical and non-ambiguous) sentence. On the current view we would have to allow that any such sentence type has a vast number of meanings, each one relating to *something* the sentence could be used to say. Such multiplying of meanings is an anathema to any kind of generalised semantic theory and would make entirely opaque how we ever learn or use a language. Secondly, however, it seems pretty clear that at least some of the above reports pick up not on what we might think of as the meaning of the antecedent sentence produced, but instead depend on features quite external to that sentence, like the conversational context in which it is produced and the (shared) background assumptions of speakers and hearers. Some of these indirect speech reports, though (potentially) perfectly acceptable, seem to be reporting things only tangentially connected with the meaning of the sentence produced. For instance, it is easy to imagine a context in which (6) is acceptable: say Jim produced (1) in an argument about the relative merits of right-wing versus left-wing Labour leaders, where it was mutual knowledge that Blair was the most right-wing leader to date. Then (6) might well be an accurate report of what Jim conveyed to his audience on this occasion by his utterance. Yet, although it is relatively easy to find a context in which (6) is acceptable, this does not seem to lead us towards taking the content sentence of (6) to tell us anything much about the meaning of the sentence Jim originally produced. (6) may well give us what Jim meant by his utterance of (1), but this seems a very different thing to what the *sentence* itself means.

The problem we are recognising here is that some perfectly legitimate indirect speech reports pick up not on sentence meaning but on what Grice termed 'utterer's meaning'.<sup>3</sup> If this is correct, then there can be no general

3 Grice 1989: 117-37.

move from facts about indirect speech reports to facts about the correct content of our semantic theory. If we group together all the ordinarily correct instances of indirect speech reports (i.e. acceptable occurrences of the ordinary language locution ‘U said that p’) then we get a rag-bag class of cases, polluted by an array of non-semantic features, which cannot play a semantically informative role. What we need, it seems, is some way to select from this nebulous set some more refined class – to rule out, at the very least, acceptable indirect speech reports, like (6) or (7), which report not on the meaning of the original sentence, but merely on what the speaker who uttered it succeeded in conveying to their audience. It is no job of a successful semantic theory to tell us *everything* which a sentence may be used to say on a given occasion of utterance, and to adopt (a) is to overlook this basic fact.

So (a) is untenable. Yet we have no reason to saddle an advocate of the standard view with anything so unappealing, for both (b) and (c) above capture the claim that there is a way to move *from* indirect speech reports *to* semantic analyses, without falling prey to the objection above. For both these more moderate views hold that the construction of a correct semantic theory for some language L will be affected by indirect speech reports, but *also* that which indirect speech reports are relevant is itself something which is constrained in some way. So, let us turn now to these more refined versions of the standard view.

## 2. A Privileged Notion of ‘What is Said’?

The problem faced by advocates of a refined, semantically relevant, notion of ‘what is said’ is how they will select from amongst the entire set of ordinarily acceptable instances of indirect speech reports just those which are genuinely semantically informative. That is to say, faced with a set like (2-7) what principle can we use for selecting the semantically relevant sub-set, the ‘what is said\*’ cases? A natural first move in answer to this question is to introduce the Gricean notion of ‘conversational implicature’.<sup>4</sup> A proposition may be conversationally implied if it can be inferred from the meaning of the sentence produced together with the context of utterance, plus knowledge of some general principles of good communication. For instance, the speaker who utters ‘Someone hasn’t done their homework again’, in a context where it is clear to all parties that the recalcitrant student is Smith, might be taken to flout the maxim of quantity (by failing to convey the amount of relevant

4 Grice 1989: 31.

information which she possesses); thus her audience may be licensed in inferring some non-literal proposition as the one actually communicated, namely that *Smith hasn't done his homework again*. For Grice, what is required to arrive at a conversational implicature is first to grasp the literal meaning of the sentence produced, then to see that this literal meaning flouts some quite general principle of good communication, and finally to infer that the literal proposition can't be the one the speaker meant to convey and therefore it should be replaced with a more suitable proposition which no longer flouts the communicative principle in question.

Relying on this inferential aspect to conversational implicatures, then, we might claim that an indirect speech report can be semantically informative just in case it reports a proposition which can provide the basis for an inference to a further proposition, but which is not itself inferentially arrived at.<sup>5</sup> Although this principle is roughly sketched, let us assume for the sake of argument that it (or something similar) does indeed hold good, providing us with a clear way to individuate speech reports picking up on conversational implicatures from those which do not. The question then is: is this enough? Is the class of acceptable indirect speech reports *sans* those which pick up on conversational implicatures a class which we can take to be semantically relevant? Consider again Jim's utterance of (1), and the small sample of acceptable indirect speech reports (2-7). Certainly some of these fit the model of conversational implicature reports, as just sketched. For instance, (6) and (7) might well be thought to arise as acceptable reports just in case Jim's original utterance, if understood literally, does not meet the requirements of some Gricean principle of good communication (say, relevance) in the context in which (1) was uttered.

However, it is not at all clear that any of the other reports fit this model of conversational implicature reporting, yet nor is it clear that we want to count all of the remaining reports as semantically informative. Take (5) for instance ('Jim said that Baby Leo's father lives in No.10'). Here the proper name 'Blair' has been replaced by the possessive 'Baby Leo's father', while the relationship 'living at No.10' has been changed to 'living in No.10'. Now we can easily imagine (5) being an acceptable report of what Jim said – say if the speaker of (5) wants to report Jim's utterance to a fellow worker in Leo's nursery, who is unaware that Leo's father is Blair. Here the case seems very different to one of straightforward conversational implicature; it is not as if the speaker of (5)

5 Though this need not be taken as a claim about the actual cognitive procedures employed by a given interlocutor; cf. the debate over the 'Availability Principle' in Recanati 1989, Nicolle and Clarke 1999, Bach 2001, and Taylor 2001.

had to work out the literal meaning of Jim's utterance, see it as somehow conversationally improper (after all, we are imagining that Jim's utterance was perfectly proper in the context in which it was uttered, it is only that, in the context of reporting, the expressions used would not be so apt), and infer her way to the content sentence of (5). Yet, on the other hand, it seems simply wrong to think that any part of the *meaning* of the *sentence* in (1) is given by the expression 'Baby Leo's father' (after all, one could perfectly well understand the sentence in (1) without any grasp of the predicate '\_\_\_ is the father of \_\_\_').

If this is right then the refined notion of 'what is said\*' in play at the moment (i.e. one which rules out only conversational implicature) is not strict enough: there are some perfectly acceptable indirect speech reports, which are not instances of conversational implicature reporting, which do not tell us about the correct semantic analysis of the original sentence. What we need is a *more* restricted version of 'what is said\*', and again the question is 'how is this to be arrived at?' At this point, it seems to me that there are two broad approaches the advocate of semantic significance for certain indirect speech acts might pursue: either our theorist may take the appropriate cases of indirect speech to be delivered through their relationship to the correct semantic theory (position (b)), or she may take the appropriate set of cases to be delivered via some non-semantic (e.g. perhaps syntactic) route (position (c)). On the former approach, the idea is roughly that lurking behind the everyday notion of 'what is said' is a semantically informed notion – something like 'literally said' – and it is this semantically informed notion which is to play the key role in construction of our semantic theory. On the latter approach, the everyday notion is thought to be redundant for semantic theorising, and instead what is introduced is an independent, technical definition of 'what is said\*', which can play precisely the semantically informative notion envisaged for it. In what follows I want to explore both these types of approach, but the claim will be, to the extent that either of them can deliver something semantically informative, the additional trip through 'what is said' is theoretically redundant.

So, let us begin with approaches of type (b), where we come to judgements of what is said *already* armed with some independent conception of semantic content. We do not expect judgements of indirect speech reports to *dictate* semantic analyses (to be made in a semantic vacuum, as it were), but instead see them as already informed by semantic evaluations. Thus advocates of (b) are unscathed by the fundamental objection to (a) above, for the crucial notion of 'what is said' is not co-extensive with the ordinary language locution (which may pick up on speaker meaning rather than literal sentence-

meaning) but is instead itself a semantically-informed notion. The important idea now is something like ‘literally said’ or ‘strictly said’, where this isolates the peculiarly *semantic* information conveyed by the utterance of a sentence.<sup>6</sup>

Now, initially at least, it seems that there is something of an air of circularity here: ‘what is said\*’ is supposed to be a semantically informative notion, yet it itself is to be semantically informed, and this may seem worrying. However, advocates of this position can dissipate the perceived problem by distinguishing between epistemic and metaphysical claims. For it may be that the special relationship between privileged indirect speech reports and the semantic theory is a matter of *metaphysics* (e.g. facts about ‘what is said’ in this privileged sense supervene directly on semantic facts), while the sense in which facts about ‘what is said\*’ are semantically informative is an *epistemic* one (e.g. they don’t, as (a) would have it, supply the raw data for our semantic theory, but can act as a guide or indicator of the contents of a correct semantic theory). In this way, judgements about what is said\* are thought to be more accessible or intuitively obvious than direct reflection on the contents of a semantic theory, thus they are available to provide an initial epistemic guide. This idea of epistemic immediacy is, for instance, endorsed by Reimer, who defines ‘what is said\*’ as ‘the meaning of the sentence uttered, relativized to a context of utterance’ (a definition she attributes originally to Grice), and suggests:

This seems to be a ‘pre-theoretic phenomenon’. For one could easily explain to a non-philosopher (of language) what is meant by Grice’s notion of ‘what is said’ and then ask her to evaluate (for truth or falsity) statements of the form: In uttering S, a speaker ‘says that’ *p*.

The intuitions to which such a person would appeal would be ‘pre-theoretical’ in the relevant sense. For (by hypothesis) our speaker is not committed to any particular theory as to the meaning (the semantic content) of the contextually relativized S...For such reasons, her intuitions are properly regarded as potential *evidence* for semantic theories.<sup>7</sup>

6 Cf. Richard 1998: 606, “A literal utterance of S literally says p just in case (according to [a correct semantic theory for English]) the semantic content of S (in the utterance) is p”.

7 Reimer 1998: 602. Recanati 1989: 309, also endorses the claim of epistemic immediacy for judgements of what is said: “I [have] made two related claims: first, that sentence meaning is something more abstract and theoretical than what is said; second, that we have ‘intuitions’ concerning what is said that serve as a starting point in the process of deter-

Now, although it seems correct to assign a degree of semantic immediacy or intuitive access to our ordinary judgements of ‘what is said’ (since competent speakers of a language have, in general, little difficulty in deciding whether to accept or reject putative reports of a given speech act), I think we must query Reimer’s suggestion that the more refined notion of ‘what is said\*’ is one to which interlocutors in general have any such intuitive or unmediated access. For consider the range of putative counterexamples to the semantic relevance of what is said, as raised by Cappelen and Lepore 1997. They point out that, to take just a selection of cases, on occasion we may licence the exchange of synonyms from the original sentence to the indirect speech report (e.g. allowing ‘bachelor’ to be reported via ‘unmarried male’), or we may allow the exchange of co-referring terms (e.g. ‘Blair’ and ‘that man’), or exchange of descriptive phrases or relational properties (e.g. ‘the Prime Minister’ for ‘the father of Baby Leo’). We may also allow reports to pick out elements of the original utterance (e.g. an utterance of ‘p and q’ reported by ‘q’), logical entailments (‘if p then q, and p’ reported with ‘A said that q’), and paraphrases of the original utterance (e.g. reports of a long fiscal statement with ‘A said that inflation will fall’).

I would suggest that, faced with a range of cases like this, we have very little reason to expect interlocutors to have *any* strong or consistent intuitions about which ones count as literal reports and which ones do not.<sup>8</sup> Furthermore, to the extent that interlocutors are able to come to judgements in this area, it seems that they must surely be *secondary* to semantic judgements – we don’t, it seems, take s and p to mean the same because we judge a report of ‘A said that p’ as a literal report of some sentence s, rather we judge a report as a literal report just in case we think s and p mean the same thing.<sup>9</sup> Consider a report of ‘that man lives at No.10’ by ‘A said that the man over there lives at

mining what the linguistic meaning of the sentence is’ (though we should note that he also denies that there is a purely semantic notion of ‘what is said’).

- 8 See Nicolle and Clarke 1999 for some experimental findings which apparently support the variability of our reporting practices, even when asked to report ‘what is said\*’.
- 9 The same sort of problem, I think, undermines Richard’s 1998 response to Cappelen and Lepore. Richard suggests that we will ultimately be able to construct a set of individually necessary and jointly sufficient conditions which yield all and only the semantically informative instances of the commonsense locution ‘what is said’. Yet, even if we agreed on the feasibility of such a project, it seems clear to me that to so much as get started on this kind of rule construction we *already* need some pretty robust *semantic intuitions*. That is to say, to decide that we will accept or reject a case as an instance of literal saying, we already need to have decided what we think the original sentence means. What drives the introduction of any of the necessary conditions is not, I would suggest, an intuition about what is *literally said* by a literal utterance of s, but simply a prior view about what s *means*.

No.10'. We will count this as a literal report only if we think that 'that man' and 'the man over there' mean the same (have the same semantic value); to assume that we can make an assessment of literal reporting prior to this kind of decision seems to 'put the cart before the horse', as it were.<sup>10</sup> Or again, take the report of (1) ('Blair lives at No.10') with (2) ('Jim said that Tony Blair lived at No.10 Downing St, London, UK'): does this count as an instance of *literal reporting*? It seems to me that this is something we simply lack any generally agreed intuitions about – in lieu of any claims about the meaning of *s* who knows whether *p* in this case counts as a literal report? The attempt to avoid circularity here fails for either we simply lack any intuitions over indirect speech reports like these, or those intuitions we can get hold of seem much better accommodated as direct semantic intuitions, concerning the meaning of the original sentence, as opposed to intuitions about some nebulous notion like 'literal reporting'.

Perhaps instead then we should look to non-semantic features to deliver the relevant set of indirect speech reports; this is the stance taken by (c) above. On this view there is still a move from appropriate indirect speech reports to the content of a semantic theory, but our earlier charge of circularity is avoided, since we are not appealing to semantic features in spelling out the notion of 'appropriateness'. Now, obviously, for this to be a substantial position, the theorist owes us an account of which non-semantic factors come into play. Yet there seems to be a good candidate here, in the form of an appeal to syntactic features. For instance, such a theorist might think to adapt something like Bach's 'syntactic correlation constraint' (SCC):

According to which what is said must correspond to 'the elements of [the sentence], their order, and their syntactic character' [*Studies in the Ways of Words*, p.87]. So if any element of the content of an utterance, i.e. of what the speaker intends to convey, does not correspond to any element of the sentence being uttered, it is not part of what is *said*.<sup>11</sup>

- 10 Reimer 1998: 602, fn.9, notes this and suggests that the key notion of 'what is said\*' might remain silent on difficult cases like replacing names with definite descriptions; but, if this is the case, then the key claim of the standard view – that there is a genuinely semantically informative notion of 'what is said' – seems to me to have been significantly weakened.
- 11 Bach 2001: 15. A somewhat similar approach can, I think, be found in Reimer, 1989: 599: "The technical aspect of Grice's notion lies in the fact that he uses 'says' in such a way that 'what is said' is *very* tightly constrained by the conventional meaning of the words uttered – considerably more so than it is in our ordinary, non-technical, sense of that locution." The problems to be raised above carry over to this approach.

Now, we should be clear that Bach himself *rejects* what I have called the standard view, for he denies that (certain) indirect speech reports can be used to dictate semantic content.<sup>12</sup> But the pressing issue from our current perspective is whether something like SCC could be of help to the advocate of the standard view. Could we use such a principle to deliver a semantically informative notion of ‘what is said\*’?

Well, I think the answer to this question really depends on how exactly the technical definition itself is to be interpreted. For the syntactic correlation constraint (as deployed by the advocate of the standard view) tells us that to fix on the privileged class of speech reports, we need to limit our attention to just those cases where the content sentence of the indirect speech report matches the original sentence in the above respects (i.e. there is a one:one correspondence between the elements of each, their order and their syntactic character). Yet spelling out what exactly SCC itself amounts to remains, it seems, less than straightforward: it requires, first, an understanding of exactly what elements count as constituents of the original sentence, and, second, it requires that the notion of ‘correlation’ be spelt out. Bach himself concentrates on the first question – determining when an element counts as a constituent of a sentence and when it is better handled as a pragmatic introduction capturing speaker meaning instead of sentence meaning. For instance, he stresses that not all genuine syntactic constituents are *voiced* constituents. Thus in the exchange:

A: Has Jim left?

B: Yes, he has.

A proper report of B’s utterance could include the verb phrase ‘left’, since this is taken to be a syntactically present but unspoken element of the sentence B produced (on the grounds that it is recoverable from the immediate linguistic context of B’s utterance). On the other hand, something like the implicit restriction to alcohol in an utterance of ‘I only drink Scotch’ is thought to be the result of pragmatic reasoning on the part of the hearer, rather than being supplied by a genuine syntactic constituent of the sentence produced.<sup>13</sup> Clearly, this issue – of when an element can be traced to a syntactic but unvoiced constituent, and when it is a novel element introduced by pragmatic processes – is one which requires a great deal of attention.<sup>14</sup> However,

12 Thus, since he wants to hold on to a semantically significant notion of ‘what is said’, he would resist the equivocation between ‘what is said’ and reported speech in this paper.

13 Bach 1981: 238.

14 See Bach 1994, Borg (manuscript), Recanati (manuscript).

the issue I really want to concentrate on here is the second one noted above, namely how should we understand the key notion of ‘correlation’? For even if we have succeeded in determining the precise syntactic constituents of the original sentence, it remains an open question what constraints the correlation relationship puts on the elements of an indirect speech act.

For instance, in some sense at least, the phrase ‘Baby Leo’s father’ in (5) *correlates to* the proper name ‘Blair’ in (1): both expressions play the role of determining an object – a person – as the subject of an ensuing predicative act, and both expressions turn out (given the way the world actually is) to deliver the very same object, but as we have already conceded (5) is not easily treated as semantically informative as to (1). Now of course an advocate of SCC may object that these two expressions – ‘Blair’ and ‘Baby Leo’s father’ – do not share a syntactic character, since the latter is (on most accounts) a quantified noun phrase and the former is (on most accounts) a referring term, but this move hardly seems to help. For what now is to stop an utterance of:

- 8) The British Prime Minister lives at No.10  
being reported with:
- 9) A said that the father of Leo lives at No.10.

Here we have two descriptive phrases in subject position (indeed two descriptive phrases containing the same number of words), so intuitively two expressions of the *same* syntactic character, which are co-extensive (in the actual world), and which stand in a correlation relationship in the sense of ‘playing the same role in the sentence’, yet still it seems we don’t want to take part of the meaning of (8) to be given by (9).<sup>15</sup>

The problem emerging is that, as yet, we can’t have an informative notion of ‘what is said\*’ because the technical terms involved in the definition of ‘what is said\*’ themselves have to be spelt out. Yet of course it is still open to the advocate of a syntactic constraint on ‘what is said\*’ to tell us more about the notion she has in mind here. For instance, it might be held that whether or not an item in the original sentence genuinely correlates to an item in the content sentence of the indirect speech report is itself a matter which can be settled by intuitions (say about whether the two items count as ‘meaning the same’); however, on this approach we run the risk of collapsing back into the

15 Furthermore, appeal to the extension of an expression across possible worlds can’t help either. Consider an utterance of ‘Blair is a cordate’ reported with ‘Jim said that Blair is a renate’.

kind of direct appeal to semantic intuitions which we rejected above. Though we might end up with an indirect speech report which coincides with the semantic analysis assigned to the original sentence, this would be entirely parasitic on prior semantic judgements, and thus the key idea that judgements about what is said can help specify the contents of a correct semantic theory would be lost. So, then, a more appealing move would be to adopt a very tight correlation relationship, one which, say, demanded *type identity* – it is not enough that there be a one:one isomorphism between elements of the content sentence of the reported speech act and elements of the original sentence, what has to get mapped in each case is *tokens of the very same type*.<sup>16</sup>

Now, one consequence of this view, I think, is that we no longer have anything which approximates to the notion we first started with, i.e. the notion of *indirect* speech. We have in effect traded in *oratio obliqua* for *oratio recta*: the only kind of speech acts which will count as semantically informative are those which actually repeat the terms used in the original utterance. More worryingly for the advocate of the standard view, however, is that it is simply no longer clear what role the additional move through ‘what is said\*’ is supposed to be playing. For the work of dictating the contents of a semantic theory is now being done by an independent principle relating syntax to semantics. What we have now is the claim that the contents of a semantic analysis for some sentence *s* is given by the *syntactic contents* of that sentence – semantic meaning is extremely tightly tied to syntactic features. But *this* is a principle we could simply apply directly to those sentences themselves. The

16 This may seem too strong a constraint, for, as Grice was well aware, certain accommodations will have to be made for indexicals and demonstratives. An utterance of ‘I am cold’ should not (always) be reported with ‘U said that I am cold’. In these cases, the syntactic character of the expressions will (help to) determine a semantic value (e.g. the speaker of ‘I’) which should then be referentially indicated in the content sentence of the indirect speech report; see Recanati 1989: 297. Bach, however, wants to argue that even this move (at least in the demonstrative case) takes us too far from the semantically relevant notion of ‘what is said’, requiring as it does, appeal to speaker intentions to determine the referent of ‘that’; see Bach 2001: 31-3. In this case, the syntactic correlation constraint would indeed require something like type identity (so that an utterance of ‘That is G’ must always be reported with ‘U said that that is G’). In personal correspondence, Bach has indicated that the notion of an interpreted logical form might be used to spell out his correlation relationship: thus, since ‘Baby Leo’s father’ contains constituents and structure that ‘Blair’ does not, ‘they do not really correlate’, i.e. they don’t share a logical form. This move may help Bach (concerned as he is with a semantically informative notion of ‘what is said’ which is divorced from indirect speech reports), but, as we will see below, such a move cannot be of assistance to any advocate of the standard view.

extra trip through ‘what is said\*’ seems theoretically otiose. The point here is not so much that such appeals to syntax fail to deliver a notion which will pick out all and only the semantically relevant instances of ‘what is said’, but rather that delivering this set now seems entirely *secondary* to claims about how the syntactic features of our language map to semantic features. Yet, if this is the case, it seems that we have failed to deliver a genuinely *semantically informative* notion. It’s the determination of the relationship between syntax and semantics that is informative, and it is merely a by-product of this that we can then isolate certain instances of the locution ‘A said that p’ where p matches the semantic content of the original sentence s.

So, the claim is that neither of the accounts of the privileged notion of ‘what is said\*’ provides us with the semantically informative notion we were after. On the one hand, there are accounts which take the set of privileged indirect speech reports to be constrained by semantic features (position (b)). Yet, to avoid circularity here, we need it to be the case that we have immediate, intuitive access to the set of semantically constrained indirect speech reports, i.e. intuitively judging whether ‘A said that the man there was happy’ counts as a literal report of ‘that man is happy’. It seems, however, that ordinary speakers are not in general able to make this kind of fine-grained judgement, and, to the extent that they are willing to make assessments of this kind, such judgements are *parasitic* on prior judgements about the *meaning* of the original sentence. Yet if this is the case, then we might as well look to these intuitions about the meaning of the original sentence as the basis for our construction of a semantic theory. On the other hand, we have accounts which seek to constrain the semantically relevant set of indirect speech reports by appeal to non-semantic features, such as syntax (position (c)). Here, however, even if the privileged, syntactic relation can be adequately spelt out, the appeal to what is said once again becomes otiose; for the work in semantic theorising is now being done by the relationship posited between syntax and semantics. Though the introduction of such a principle *allows* us to arrive at a set of (direct or indirect) speech reports which coincide with the content ascribed to a sentence by our semantic theory, the whole process is once again entirely secondary to the fundamental processes involved in construction of that semantic theory.

The path to delivering a semantically informative notion of ‘what is said’ seems, then, fraught with difficulties; so why not simply reject the claim that indirect speech reports are in any way semantically relevant, i.e. why not reject the semantic significance of ‘what is said’? Well, one reason to resist this kind of move might be the thought that we need to preserve some kind of intimate connection between the two domains – indirect speech reports must

have *something* to do with the content of a semantic theory. That is to say, it might seem initially that the only alternative to the standard view is a position like (e) above, whereby there is thought to be *no* relationship between indirect speech reports and a semantic theory. In the final section I want to argue that, were this the case, it might well throw us back to the so far fruitless search for a way to delimit a semantically significant notion of ‘what is said\*’, but this is not the case, for opponents of the semantic significance of indirect speech reports have another option open to them.

### 3. The Proper Connection Between a Semantic Theory and Indirect Speech Reports

To reject the standard view we must hold that there is no route *from* indirect speech reports *to* the content of a semantic theory, i.e. we must reject positions (a-c). Now, of course, one way to deny this claim is simply to deny that there is any relationship between the two realms at all; this is the claim made by (e). On (e) indirect speech reports and semantic theories belong to different domains of inquiry and each can proceed without any regard for the other. Yet it seems clear that (e) is not a very attractive position – it might not be logically impossible that there be a correct theory of meaning for some language L which plays no role at all in specifying the content of communicative acts involving the use of L (perhaps if L was a language used by a race who had direct, unmediated access to one another’s communicative intentions, so that language played for them some other, non-communicative role), but this certainly doesn’t seem plausible for our own cognitive condition. We simply cannot allow that there is no connection at all between what a correct semantic theory states for a given sentence and (correct) judgments of what that sentence may be used to say. It is perhaps the desire to avoid such an obviously unattractive position as (e) that motivates the search for a semantically relevant notion of ‘what is said\*’. However, as has been evident since §1, (e) is not the only position available to those who want to deny the standard view. Instead they might pursue position (d), whereby there is a one-way connection between the two realms, but it runs only *from* semantic analyses *to* indirect speech reports. It is this position which I wish to advocate.

The claim I want to make is that the output of a semantic theory is relevant to determining acceptable or correct indirect speech reports, but that this is just because the output of such a theory provides the starting point for working out what given speakers communicate by an utterance of a given sen-

tence. Since the semantic analysis of a sentence provides the initial input to constructions of indirect speech reports, these latter reports are clearly dependent on the output of a semantic theory, but what is important to note is that this does *not* entail that there is any clear or accessible route back from judgements about what is said to the output of a semantic theory. What precludes *this* is that, on the model I would like to promote, by the time we reach any judgement of what is said by a given utterance in a given context the peculiarly semantic contribution will have been swamped beneath the influence of a vast range of other (non-semantic) factors. Finding out what is said is always, I would argue, finding out what someone said, sometime, someplace; it is, at heart, a notion which subsumes much more than the mere semantic, and this is why we should not expect a semantically informative notion to emerge from such essentially pragmatic considerations.

To sketch the kind of picture I have in mind, then, let me borrow from the framework made familiar by Chomsky, Fodor and others. I assume a broadly modular picture of the mind, containing discrete bodies of information, and encapsulated processes acting on that information, dealing with such subjects as perception, commonsense psychology (or theory of mind), commonsense physics, and, of course, language. The language faculty, as I conceive it, is comprised of at least three sub-domains: orthography and vocalised speech recognition, syntax and semantics. The semantic information the faculty contains is of quite a limited form; say just that required for generating the truth-conditions of complex linguistic items on the basis of the meaning of their parts and their mode of composition. It is this quite constrained item (the ‘minimal proposition’ in Recanati’s terminology) which feeds out of the language faculty and *at this point* becomes subject to a wealth of other information which the agent possesses. Judgements of what is said, I want to suggest, are *essentially* the output of this further process – what we get *only* once the purely semantic contribution has been subjected to a range of other (non-semantic) information the speaker takes to be relevant in the situation. Thus we cannot, contrary to the standard view start with indirect speech reports and move from here to the semantic content of sentences. It is, I suggest, the task of a successful semantic theory to account for certain semantic facts (such as the productivity and systematicity of our language) and not to account for the vast range of factors which are involved in making assessments of what is said. The proper domain of semantics is the *meaning* of simple and complex linguistic items, and shifting our focus from this to the tangential relation of ‘what is said’ serves only to confuse matters – indirect speech reports supervene on the output of a correct semantic theory, but

it is a mistake to think such speech reports can reveal anything about the proper content of a semantic theory.

Theorists like Recanati are, then, I believe, right when they suggest that pragmatic processes are necessary prior to any analysis of what is said.<sup>17</sup> Thus, as Recanati writes:

There is a single notion of what is said, and that is a pragmatic notion: saying, as Grice claimed, is a variety of non-natural meaning, characterised by the role which the conventional meaning of the sentence plays in the hearer's intended recognition of the speaker's communicative intentions ... There is ... no purely semantic notion of what is said.<sup>18</sup>

This recognition is at odds with the claims of a formal semantic theory (i.e. that we can deliver accounts of sentence meaning in a context-independent manner) *only if* 'what is said' is thought to be semantically relevant. Yet, as I have tried to argue, we have no reason to think this is the case. The facts about indirect speech reporting do not, from the start, support the idea that 'what is said' is a semantically significant notion (as Cappelen and Lepore have shown). Furthermore, if we try, in the light of this, to replace the ordinary locution with some technical notion, 'what is said\*' which delimits all and only the semantically relevant cases, things do not seem to improve at all. For it seems the only option for yielding this refined notion of 'what is said' is to appeal either to semantic properties or to syntactic properties. However, as I have argued, the first move is unable to avoid the charge of circularity, while the second approach makes the appeal to 'what is said' theoretically otiose. Thus we simply cannot deliver the technical notion necessary to support claims of semantic significance for indirect speech reports, but nor, I have suggested in this section, should this finding worry us. For so long as we can preserve the connection from semantic analyses to indirect speech reports we avoid collapse into the unpalatable position (e), yet without embracing the idea that indirect speech reports are semantically informative, i.e. without holding that there is a semantically significant notion of 'what is said'.<sup>19</sup>

17 Where I think Recanati and other advocates of what he labels 'truth conditional pragmatics' (e.g. Sperber and Wilson's 1986 Relevance Theory), go wrong is in thinking that this shows that the correct form for a theory of linguistic meaning must be one which is pragmatic 'all the way down' as it were. The argument against truth-conditional pragmatics is explored in more length in Borg (manuscript).

18 Recanati 2001: 87-8.

19 Thanks are due to Kent Bach and Ernie Lepore. Research for this paper was made possible by the award of a Mind Fellowship.

## Bibliography

- Bach, K. "Referential/Attributive." *Synthese* 49 (1981): 219-44.
- . "Semantic Slack: what is said and more." In *Foundations of Speech Act Theory*, ed. S. Tsohatzidis. London: Routledge, 1994: 267-91.
- . "You Don't Say?" *Synthese* 128 (2001): 15-44.
- Borg, E. "Saying What You Mean: Unarticulated constituents and communication" (manuscript).
- Cappelen, H. and E. Lepore. "On an Alleged Connection Between Indirect Speech and the Theory of Meaning." *Mind and Language* 12 (1997): 278-96.
- . "Reply to Richard and Reimer." *Mind and Language* 13 (1998): 617-21.
- Grice, P. *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press, 1989.
- Nicolle, S. and B. Clark. "Experimental Pragmatics and What is Said: a response to Gibbs and Moise." *Cognition* 69 (1999): 337-354.
- Recanati, F. "The Pragmatics of What is Said." *Mind and Language* 4 (1989): 294-328.
- . "What is Said", *Synthese* 128 (2001): 75-91.
- Recanati, F. "Unarticulated Constituents," (manuscript).
- Reimer, M. "What is Meant by 'What is Said'? A Reply to Cappelen and Lepore." *Mind and Language* 13 (1998): 598-604.
- Richard, M. "Semantic Theory and Indirect Speech." *Mind and Language* 13 (1998): 605-16.
- Sperber, D. and D. Wilson. *Relevance*, Oxford: Blackwell, 1986.