

This is a preprint of a Paper accepted for publication in the *Proceedings of the Aristotelian Society* © 2009

Semantic Compositionality and Truth-Conditional Content

Alison Hall

It is widely held that hearers grasp an utterance's truth conditions by assigning contents to the linguistic expressions used, and combining these contents according to semantic composition rules. To preserve compositionality of truth-conditional content while accounting for context-sensitivity that is not traceable to *overt* linguistic form, semanticists posit *covert* linguistic structure. The strongest justification for this approach is the allegedly unconstrained nature of the alternative, whereby a process of 'free pragmatic enrichment' supplies constituents of content that are not traceable to (overt or covert) encoded meaning. This paper argues that free enrichment is tightly constrained by purely pragmatic factors, thus undermining the motivation for semantic compositionality.

I. Introduction

It's generally agreed that the truth-conditional content of an utterance can go well beyond what the overt (pronounced) material seems to provide. In each of (1)-(3), (a) is the sentence uttered, while (b) is a (very rough) indication of the possible truth-conditional content, in an appropriate context¹:

- (1) a. It's raining.
b. IT'S RAINING IN LONDON
- (2) a. Every student passed the exam.
b. EVERY STUDENT IN MY CLASS PASSED THE EXAM
- (3) a. Fixing the car will take time.
b. FIXING THE CAR WILL TAKE A LONGER TIME THAN EXPECTED

Pragmatists (including Carston 1988, 2002; Recanati 2004; Sperber and Wilson 1986/1995) use such data to argue for 'unarticulated constituents' of truth conditions, not traceable to any element of standing linguistic expression-type meaning. Instead, these constituents are provided entirely on pragmatic grounds by a process of free (i.e. not linguistically mandated) enrichment, which involves the pragmatically motivated and controlled development of a subpart of the standing meaning. Pragmatic processes generate these constituents to meet expectations of relevance, informativeness, and so on: for example, (3) is trivially true without the enrichment – every activity takes some amount of time or other, however minimal – and so a proposition that resulted from just decoding plus reference assignment would have no non-trivial implications. With (1), the information that it's raining (*punkt*, or somewhere) would be little use to hearers, and would not be something that the speaker intends to communicate in its own right; to derive further information from the utterance, the hearer needs to infer from the context the location of the rain. In all

¹ Small capitals represent propositional contents.

these cases, the result of just decoding plus linguistically mandated pragmatic processes (reference assignment, disambiguation) is not a proposition that would be speaker-meant, and is not the content that provides the basis for inferring the implicatures of the utterance; for this, the freely pragmatically supplied elements are required.

Many authors deny the existence of free enrichment. Stanley (2000, 2002), Stanley and Szabó (2000), and King and Stanley (2005), among others, accept that truth-conditional content exceeds overt material, yet maintain that ‘all truth-conditional effects of extra-linguistic context can be traced to logical form’ (Stanley 2000: 391). Their reaction to examples such as those above is to posit covert indexicals encoded in the linguistic expression-type meaning. For example, the verb ‘rain’² has an attached location variable, which is assigned an appropriate value in context, as illustrated in (1). In (2), the quantifier domain is restricted by assigning a value to a domain variable encoded with the noun ‘student’ (or with the quantifier ‘every’, on some semantic accounts of domain restriction). Motivating this ‘semanticist’ approach is the desire to preserve a systematic, compositional explanation of our understanding of truth conditions: it succeeds because we know what the words used refer to, and understand how their contents are combined.

The case for the semanticist’s covert variables is based largely on interpretive, rather than linguistic, considerations (see, e.g. King and Stanley 2005; Szabó 2001)³, leaving him in the uncomfortable position of positing extensive syntactic structure for which there is no syntactic evidence. From the point of view of linguistic theory, the free enrichment approach would have a clear advantage: leaving more responsibility to pragmatics makes for a simpler, more elegant syntax and semantics, while the pragmatic mechanisms and principles involved in free enrichment are just those that are independently needed for other pragmatic tasks such as implicature calculation. Furthermore, it has been argued that compositionality of standing linguistic meaning, rather than of truth-conditional content, satisfactorily explains how our finite minds are capable of producing and interpreting infinitely many novel sentences, and why lexical items make systematic contributions to the meanings of the various sentences in which they appear (see, e.g., Powell 2002; Carston 2002, p. 70-4). The motivation for covert indexicals relies crucially, then, on the assumption that truth-conditional content must be semantically compositional, as the alternative is intractable: Free enrichment is claimed to overgenerate interpretations of sentences; it appears too powerful and unconstrained.

The ‘overgeneration’ objection is that, while pragmatic contributions to truth conditions that do occur can be accounted for in terms of pragmatic principles, pragmatic theories don’t seem to make clear predictions about where free enrichment *can’t* take place (this objection is presented forcefully in Stanley 2002). (The hidden indexical account, in contrast, has, in principle, a simple explanation: no pragmatic

² Underlining indicates that a linguistic expression is being mentioned.

³ Stanley (2000) attempts to provide linguistic evidence (from binding) for certain covert variables. As this has already been much disputed (cf. Carston 2002; Neale 2004; Recanati 2004; Hall 2008b), I do not address it here.

effects on truth conditions occur where not mandated by the linguistic form⁴.) Perhaps the most serious concern about the pragmatic account – and the one I address here – is that it seems to allow that extra propositions, or extra semantic arguments or predicates (such as NP- or VP-conjuncts), could be incorporated into truth-conditional content; however, it is agreed that they can't be. As Stanley (2002) and Elbourne (2008) point out, though, it looks stipulative to exclude them on an account that relies on a powerful pragmatic inferential capacity able to freely 'intrude' on truth conditions, supplying unarticulated constituents such as those in (1)-(3) without linguistic mandate. In the next section, I respond to this charge and show how the incorporation of inferred propositions, semantic arguments and predicates into the utterance's truth-conditional content is excluded by general considerations about pragmatic processing and rationality, without the need for any extra constraints from linguistic meaning.

II. A response to the overgeneration charge

It is clear that, as a matter of empirical fact, extra, inferred semantic arguments, predicates, and propositions, cannot form part of the utterance's truth-conditional content. However, pragmatists see the truth-conditional content as a development of the linguistic logical form, so the constituents of logical form are preserved. One might ask, then, why extra conjuncts (for example) cannot be added, as long as the original logical form is preserved within the 'enriched' proposition. An example would be the quite common type of case where the truth-conditional content of an utterance is 'visible' in one of the implicatures, as happens in one of the cases of alleged overgeneration suggested by Stanley. Imagine that the contextual assumption (4) is already highly salient. In that case, Stanley asks, why can't an utterance of sentence (5) be enriched to the proposition (6)?

(4) EVERYONE_i WHO LIKES SALLY LIKES HIS_i MOTHER

(5) Everyone likes Sally.

(6) EVERYONE_i LIKES SALLY AND HIS_i MOTHER (Stanley 2002: 165-6)

The answer, in this case, is quite straightforward: The proposition EVERYONE LIKES SALLY (setting aside domain restriction) is needed independently as input to a modus ponens inference together with the premise in (4) to derive the proposition EVERYONE LIKES HIS MOTHER, before the two could, even in principle, be conjoined. And EVERYONE LIKES SALLY is not a contextual assumption, in the context described, so must be derived as the truth-conditional content of (5): the development of the logical form into truth-conditional content cannot go beyond this and add extra constituents, because this proposition is required independently as input to warrant further inference – to implicatures or implications.

The explanation for that example can be generalized to a range of cases: if an assumption (developed from the logical form) is needed as a premise in the derivation of further intended aspects of meaning, as is frequently the case, then it cannot be

⁴ In practice, though, the challenge for the semanticist is to show that *every* pragmatic effect on truth conditions is linguistically mandated. If he can't, then he is open to the same overgeneration objection that he levels at the pragmatist, and the motivation for the semanticist account is undermined.

developed any further at the level of truth-conditional content by incorporating extra propositions or conjuncts, etc. Further development would block the inference, and thus prevent any inferential warrant for obviously intended conclusions. Any conjoining of it with such additional material must therefore take place at a different stage, namely, implicature derivation. In Hall (2008a) and (2008b), I apply this argument to a range of other specific examples; here, though, I discuss the more general explanation for why this kind of ‘enrichment’ does not occur, which has to do with the way in which the overall interpretation is warranted.

The key point about the distinction between truth-conditional content and implicatures, on the approach I am defending, is that truth-conditional content is a development of the logical form of the utterance, so is derived by a combination of decoding and pragmatic processes, while implicatures are derived purely inferentially from a set of fully propositional premises. Because implicatures are inferred as conclusions from the premises consisting of truth-conditional content and contextual assumptions, the overall interpretation forms a valid argument, with the truth-conditional content and contextual assumptions warranting the implicatures⁵. That the overall interpretation that a hearer recovers forms this valid argument is what makes the process inferential, as opposed to being, for example, a purely associative or connectionist process.

Although logical form is *logically* prior to truth-conditional content, and truth-conditional content to implicature, most authors (including Recanati 2004, p. 49-50; Stanley and Szabó 2000, p. 230-1; Wilson and Sperber 2002) agree that online comprehension cannot simply be, and in fact, generally is not, a matter of first decoding the logical form, then recovering the truth conditions, and only then calculating implicatures. Rather, the hearer’s hypotheses about intended implicatures and implications, formed on the basis of his expectation of relevance given the conversational situation, can influence the development of logical form into truth-conditional content, and his retrieval or construction of contextual assumptions. As Wilson and Sperber (2002) put it, the comprehension process involves hypotheses about truth-conditional content, implicatures, and contextual assumptions being ‘mutually adjusted’, in parallel, until the various assumptions settle into a valid argument relation with truth-conditional content and contextual assumptions warranting implicated conclusions (and this overall interpretation meets the hearer’s expectation of relevance). This process can involve several adjustments and readjustments to each of the various kinds of assumptions involved, with hypotheses about any one or combination of truth-conditional content/implicature/contextual assumption affecting hypotheses about any of the others⁶.

⁵ An inference is sound, or warranted, if it is based on premises from which it follows deductively. For example, given the premises in (i) and (ii), the conclusion in (iii) follows deductively and is warranted, in as strong a sense of those expressions as can reasonably be expected to apply to pragmatic inference:

- (i) If it’s sunny, we can play tennis.
- (ii) It’s sunny.
- (iii) We can play tennis.

⁶ The following example provides a simple illustration of mutual adjustment:

- (i) A: Do you want to come round for dinner tonight?
B: I’m going to the cinema.

So truth-conditional content is derived by both decoding and pragmatic processes; implicatures and contextual assumptions are entirely pragmatically inferred. With linguistic decoding, the hearer can normally be certain that the meaning he recovers is something that the speaker intended him to use in working out her meaning, as decoding is an algorithmic, invariant process. Pragmatic inference, in contrast, always involves some risk of misunderstanding, because it depends on the hearer being able to figure out the speaker's intentions, and on the speaker judging correctly what the hearer can figure out, what interpretations are accessible to him, and so on. Inferential comprehension is a matter of forming and confirming a hypothesis about the set of assumptions the speaker intends to communicate, and the confirmation of the interpretation is constrained by the need for the various propositions communicated by the utterance to form a valid argument relation, with the premises (contextual assumptions, truth-conditional content) warranting the conclusions (implicatures). Because, in ostensive communication, a speaker is trying to get her message across, and can to some extent predict what information is accessible to the hearer and what interpretations are likely to occur to him, the first interpretive hypothesis to occur to the hearer has a high degree of initial plausibility, simply by virtue of occurring first (see Sperber and Wilson 1986/1995 and Carston 2002 for detailed justification of this claim). So on the hearer's part, his initial hypothesis – say a hypothesis about an intended implicature – is given initial warrant by his expectation of relevance, and if, by developing the logical form and accessing contextual assumptions, he can form an argument on which this implicature is warranted by premises that are accessible to him (premises whose accessibility to him the speaker should have been able to predict, and which the utterance guides him to), then this increases the likelihood that the overall interpretation is the intended one.

Hearers have a strong warrant for the overall interpretation, then, if they can get the various assumptions involved to form this relation where there is a sound inference from the truth-conditional content and contextual assumptions to the implicature. And, as I indicated above, what makes the process inferential is that the various stages are all constrained by the fact that the overall aim is this sound inferential relation.

Implicatures are properly inferentially warranted, because they follow logically from the premises. Between logical form and truth-conditional content, however, there is no relation of (deductively) valid inference, and free enrichments, being subpropositional constituents, do not follow logically from anything; rather, they are recovered on the basis of their high accessibility in the context of utterance (and confirmed in so far as they contribute to an overall interpretation which is optimally relevant). This excludes the addition of inferred propositions/semantic arguments/predicates to truth-conditional content for the following reason.

The idea is that with B's reply here, the hearer, A, would first form a hypothesis about the implicature, from various cues, including perhaps B's facial expression, the fact that she's starting to explain and offer justification, which suggests that she's communicating a negative answer to A's question. This negative answer is an implicature, but it seems plausible that it could be accessed first, and only subsequently would the temporal reference be fixed (to TONIGHT) in the truth-conditional content, and the contextual assumption constructed (something along the lines of going to the cinema precluding going to dinner with A on the same evening): truth-conditional content and contextual premise are being adjusted to warrant the implicature.

While the implicature is warranted by following logically from the premises (the truth-conditional content and contextual assumptions), the contextual premises selected and the truth-conditional content derived receive their inferential warrant in different ways. Given mutual adjustment, described above, the premises can be confirmed by ‘backwards’ inference: that is, if the conclusion – the implicature – seems a promising hypothesis about the speaker’s intended meaning (for instance, it would answer the hearer’s question), and the entire interpretation is confirmed by being consistent with expectations of relevance, then the hearer has good reason to adjust the premises so that they warrant that conclusion. Contextual premises get their warrant entirely from this kind of confirmation – basically, they are warranted by the fact that they fit into this valid argument relation with the other assumptions that are in play – and aren’t constrained by any logically prior stage in the interpretation process. The truth-conditional content, on the other hand, while needing to form part of this argument relation, also needs to be justified given the (logically prior) linguistic meaning – if it were just warranted by fitting in to the argument relation, the inferential link between linguistic meaning and truth-conditional content would be lost (it would be essentially just an associative relation). The move from linguistic meaning to truth-conditional content needs to be as rational a step as possible, given that it is part of an overall inferential process that should be reconstructable by explicit reasoning. But because this particular sub-part of the overall process is not logically warranted, that is a good reason for the material that is composed into truth-conditional content to be, in a sense, minimal. Other propositions, or semantic arguments/predicates such as NP- or VP-conjuncts (which can be straightforwardly propositionalized), that are not partially isomorphic with the linguistic meaning, can stand alone, and therefore will remain as independent propositions, as this way they can be inferentially warranted by virtue of forming an independent premise or conclusion in the argument that is constructed as the interpretation of the utterance. Subpropositional constituents, such as those that are composed into the truth-conditional content as unarticulated constituents in (1)-(3) and the other examples of enrichment that pragmatists (e.g. Carston 2002; Recanati 2004) have discussed, cannot form independent premises/conclusions, so cannot themselves be warranted in the same way. So, as long as there is enough evidence of a different sort for them – for instance, they are highly accessible in the context of utterance and are compatible with the linguistic meaning – they are incorporated into the truth-conditional content. Free pragmatic processes affecting truth-conditional content, then, are constrained to be *local* processes, rather than global inference, and the pragmatist’s account is not susceptible to the accusation that it overgenerates by allowing for extra inferred conjuncts, and so on, to ‘intrude’ on truth-conditional content.

III. Conclusion: free enrichment as a local process

Free pragmatic enrichment is a process of local modifications of subparts of the encoded logical form, and, as Recanati (2004) puts it, it is the modified meaning of these subparts, which can contain elements not traceable to the linguistic semantics, that goes into the composition process. Some degree of enrichment is often required, but it occurs locally, and just as far as is necessary to reach a proposition that provides the inferential warrant for the implicatures and implications that the hearer draws from the utterance. This can be illustrated by the following examples (material outside

brackets is the uttered sentence; inside brackets are possible unarticulated constituents):

- (7) The ham sandwich [ORDERER] wants his bill.
- (8) It will take [A LONG] time.
- (9) Every boy [IN THE CLASS] was there.
- (10) I've got nothing [SUITABLE FOR A WEDDING] to wear.

Recanati (1993, 2004), Sag (1981) and Nunberg (1995) discuss metonymies such as (7), and point out that (when it's appropriately contextualised) we don't seem to first compute the absurd 'literal' meaning on which a culinary item wants the bill, then, recognizing the absurdity, infer that the speaker was referring to the person who ordered it. Instead, the deferred meaning is computed at the local level, and is what goes into the composition process. It is also clear why the enrichment is necessary here: for example, the speaker could be implicating that the hearer should give the customer the bill, and for this inference to go through, the expression 'ham sandwich' needs to be taken to refer to the customer. Enrichment of (8) involves modification of just the noun, rather than recovery first of the trivially true proposition that the activity in question will take place over a period of time, and then the calculation that what the speaker is trying to communicate is something else. Likewise, domain restriction, as in utterances of (9) and (10), can also be seen to be local: Recanati (1993, p. 262-3) treats (9) as enrichment of the encoded predicate BOY to BOY IN THE CLASS, rather than the whole encoded propositional form being enriched.

Free pragmatic enrichment is, then, a local process. Assuming that pragmatic inference is conducted not in natural language, but in a distinct, modality-neutral, conceptual medium (a Mentalese or Language of Thought, as defended in Fodor 1975), then pragmatic enrichment can be thought of as roughly equivalent, at the language-of-thought level, to the addition of linguistic adjuncts in speech (elements that modify a subpart of the encoded form, rather than all of it). The addition of semantic arguments/predicates or propositions, in contrast, is not a case of just modifying a *subpart* of the linguistically encoded meaning; rather, it is a global process in that it operates on fully propositional forms (cf. Recanati 2004, chapter 2).

Once enrichment is seen to be a local pragmatic process, it becomes far more theoretically tractable. In principle, *any* local enrichments may be possible, but whether or not they occur will depend partly on factors such as accessibility, and compatibility with expectations of relevance, that apply to pragmatic processes generally – for example, the amount of enrichment will be constrained by the fact that pragmatic processing will take place only as far as it will yield effects for no gratuitous effort (an idea defended in Sperber and Wilson 1986/1995). In Hall (2008b), I suggest that, from the distinction between local and global pragmatic processes that I've discussed here, there emerges a further constraint on unarticulated constituents: as well as being local, they cannot consist of information that is 'at-issue' in the context, because such information, to be accepted as part of the interpretation, needs to be given the warrant of a logical argument. This latter idea remains to be worked out in detail, but together, these two constraints should allow for the unarticulated constituents that do occur, and preclude those that don't, thus

removing the motivation for semantic compositionality as an explanation of how we grasp an utterance's truth-conditions⁷.

Department of Linguistics
University College London

UCL Linguistics
Chandler House
2 Wakefield Street
London WC1N 2PF
U.K.
a.hall@ucl.ac.uk

References

- Carston, R. 1988. Implicature, explicature and truth-theoretic semantics. In Kempson, R. (ed.), *Mental Representations: The Interface between Language and Reality*, 155-81. Cambridge: Cambridge University Press.
- 2002. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Elbourne, P. 2008. The argument from binding. To appear in *Philosophical Perspectives*.
- Fodor, J. 1975. *The Language of Thought*. New York: Thomas Crowell.
- Hall, A. 2008a. Free enrichment or hidden indexicals? *Mind and Language* 23: 426-56.
- 2008b. *Free Pragmatic Processes and Explicit Utterance Content*. University of London PhD dissertation.
- King, J. and Stanley, J. 2005. Semantics, pragmatics, and the role of semantic content. In Szabo, Z. (ed.) *Semantics vs. Pragmatics*, 111-64. Oxford: Oxford University Press.
- Nunberg, G. 1995. Transfers of Meaning. *Journal of Semantics* 12: 109-32.
- Powell, G. 2002. Underdetermination and the principles of semantic theory. *Proceedings of the Aristotelian Society* 102: 271-8.
- Recanati, F. 1993. *Direct Reference: From Language to Thought*. Oxford: Blackwell.
- 2004. *Literal Meaning*. Cambridge: Cambridge University Press.
- Sag, I. 1981. Formal semantics and extra-linguistic context. In Cole, P. (ed.) *Syntax and Semantics 13. Radical Pragmatics*, 273-94. New York: Academic Press.
- Sperber, D. and Wilson, D. 1986/1995. *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Stanley, J. 2000. Context and Logical Form. *Linguistics and Philosophy* 23: 391-434.
- 2002. Making it articulated. *Mind and Language* 17 1&2: 149-68.
- Stanley, J. and Szabó, Z. G. 2000. On quantifier domain restriction. *Mind and Language* 15: 219-61.

⁷ I'm grateful to Robyn Carston for many discussions of and comments on the material in my PhD dissertation, from which this paper is drawn. This work was supported by an AHRC doctoral award.

- Szabó, Z. 2001. Adjectives in context. In Kenesei, I. and Harnish, R. M. (eds.) *Perspectives on Semantics, Pragmatics, and Discourse*, 119-46. Amsterdam: John Benjamins.
- Wilson, D. and Sperber, D. 2002. Truthfulness and relevance. *Mind* 111: 583-632.