Reply to Hawthorne

In Chapter 7 of *IS* we rely crucially on tests for how speakers share content across contexts. We claim these tests can be used to gather evidence both for and against claims about an expression being context sensitive. Many philosophers now rely on these and related tests – Hawthorne (2003) being early proponent (cf. also Egan, Hawthorne and Weatherson (2004), Lasersohn (2006), Macfarlane (2004), Richard (2004), and (arguably) Stanley (2005)). In his reply, Hawthorne raises interesting challenges to our use of such tests. We agree that the issues are important and that thinking them through will help clarify the nature of the evidence they provide. But we disagree with him about whether the challenges he raises threaten anything we assert in *IS*. We see Hawthorne's comments as one more stage in an active research program of which *IS* is a part.

Hawthorne's Examples: 'left' and 'nearby'

Hawthorne uses the expressions 'left' and 'nearby' to test our indirect disquotational reporting test. According to this test, if an expression e is context sensitive, it should not be possible to indirectly report a particular use of e in a context C in arbitrary contexts C' by disquoting it. The first person pronoun 'I' is a bona fide context sensitive word and as a matter of fact we cannot accurately and arbitrarily disquote its use in any context. More specifically, we are unlicensed to report an utterance of 'I am F' by a speaker S in C in an arbitrarily different C' with an utterance of 'S said that I am F'. [1] We hope this is obvious. On this test Hawthorne comments:

> Suppose Ernie is in New York City and I am in Birmingham. Ernie says 'A nearby restaurant has good Vietnamese food'. I can report this by saying 'Ernie said that a nearby restaurant has good Vietnamese food,' even though I am far away from him. Suppose that Ernie is facing me. A car goes to Ernie's left and my right. Ernie says 'The car went left'. I can say 'Ernie said that the car went left', even though my orientation is radically different to his

Assuming 'left' and 'nearby' are context sensitive, with distinct utterances picking out distinct directions and locations, if the data are accurate, we are left with a choice, Hawthorne says. We can either relinquish the test or endorse the view that 'left' and

[1] See Chapter 7 for further elaboration on these tests.

'nearby' are semantically invariant. The latter option, says Hawthorne, is 'horrible', so the test should go.

Five comments on Hawthorne's discussion of these cases:

1. 'Left', 'Nearby' and note 1 of IS: On the first page of *IS* we mention so-called contextuals with 'left' and 'nearby' being two. We add them to Kaplan's (1989) Basic Set of context sensitive expressions (i.e. we do not advocate invariantism about them), but in the first footnote of *IS* we say: 'To be honest, we have our doubts about so-called contextuals; and it's probably no accident that they did not occur on Kaplan's (1989) original list. We will let you decide for yourself after you have read our book.' When we wrote *IS* we were, quite frankly, befuddled by contextuals. Our sense was that the data about their usage were messy. Cleaning up the mess was not central to *IS*, and so, we never returned to the issue. Hawthorne's reply has clarified the source of our confusion. Before turning to it (in §§4-5 below), some comments on the data and Hawthorne's claim that the data should lead us to 'give up' our tests are in order.

*2. Why 'giving up tests' is not an Option***:** Hawthorne thinks these examples should lead us to 'give up the tests'. First a brief comment on why this is not an option. We never presented our test as necessary or sufficient for context sensitivity. Rather, we describe them as providing 'evidence' that expressions are context sensitive (see, e.g., (p. 89), where we introduce the tests);[2] ´the challenge', we say, 'is to explain away or challenge this data (or show that the test is no good)' (p. 96; also p. 99).[3] The test is, if you like, a procedure for collecting data – data that has got to be explained (away) by *any* theory. Applied to 'left' and 'nearby' our claim is this: if you can disquotationally report utterances of 'left' and 'nearby' in relevantly different contexts (for more on 'relevantly different contexts' see below) and you *don't* want to claim they are context sensitive, then you either have to explain the data away (i.e. present a theory that accounts for them without holding they are semantically stable) or challenge the data (i.e. claim that the intuitions are faulty in some way). This is exactly how Hawthorne proceeds. He presents two explanations – exactly what we claim should happen in such circumstances. His data couldn't lead us to 'give up' the test. On the contrary: they strengthen our resolve in it.

---

[2] Pagination is to *Insensitive Semantics* (2005) unless otherwise noted.
[3] For further discussion of this see our 'Reply to Critics' in *Mind and Language* (2006).

3. Diagnosis of Hawthorne's Examples First, notice: we can disquotationally report uses of 'I' *if the right circumstances are in place*, i.e. if the reporter is identical to the original speaker. For the test to be applicable, it is essential that the contexts of the report and the reportee be *relevantly different*, i.e. that the relevant contextual features should be different in the contexts of the report and the reportee.

So what are the contextually relevant features for, say, 'nearby'? Maybe the answer is something like this: the character of 'nearby' determines that an utterance of 'nearby' in a context C refers to a location salient in C, but that location needn't be the location of C (i.e. it needn't be the location where the speech act occurs). It is whatever location is salient in C, and that could be the location of the reportee – if, for example, that location is made salient in indirectly reporting her.[4]

Hawthorne's examples can then be understood as examples in which we can disquotationally report utterances containing 'nearby' because the relevant contextual features in the two contexts (i.e. the salient location) are the same. If this is the correct diagnosis, all the examples show is that when the same location is salient in two context, C and C', you'll be able to disquotationally report utterances of sentences containing 'nearby' from C to C' (or the other way around). But this is no more surprising than being able to disquotationally report utterances of sentences containing the first person pronoun when the reporter is the same as the reported speaker.

According to this suggestion, 'nearby' refers to whatever location is salient in the context of utterance, and if your intuition is that Hawthorne's disquotational reports are true, that is because the location of the reportee has become salient in the context of the report. Further support for this diagnosis is provided by imagining examples in which the salient locations differ in the context of the report and the context of the reportee, and hence blocks the disquotational report. Consider this variation in Hawthorne's example (i)-(ii):

i. Ernie says to John, walking on 7th street in NY looking for a restaurant: *A nearby restaurant has good Vietnamese food.*

[4] This is close to a possibility Hawthorne considers when he says: ''And there is an easy explanation of this: 'Ernie is nearby' is true relative to a contextually relevant location that is supplied by context, but the *relevant location need not have anything to do with the location where the utterance is made*.''

ii. John, walking around Birmingham looking for a restaurant, reports Ernie's utterance to his friends by saying: *Ernie says that a nearby restaurant has good Vietnamese food.*

(ii) is intuitively false, and the above account provides an explanation: in the context of (ii) it is the location of the report (i.e. the location where the indirect report is uttered) that is salient - the search for a restaurant in Birmingham has made Birmingham the salient location.

5. Mixed Quotation: Hawthorne says that ‘‘many 'say that' reports have a feel of 'mixed quotation'’’ – where only a part of the report directly quotes (and so mentions) a part of the utterance reported. Applied to 'left' and 'nearby', his idea is that they are within non-articulated quotation marks, even though the rest of the report is indirectly quoted (and therefore, used). Properly understood his examples should be rewritten as:

(1) Ernie said that a ‘nearby’ restaurant has good Vietnamese food.

(2) Ernie said that the car went ‘left’.

Hawthorne concludes that ‘the first test is a terrible test for context-dependence, unless we can screen off 'mixed quote' readings’ – because the test works only if we can ensure we are dealing with genuine indirect reports. Call the requirement that we use the indirect reporting tests only if we know how to distinguish mixed from indirect quotation 'HR' (for Hawthorne's Requirement.)

Three comments:

a. Hawthorne tries to motivate HR by an appeal to examples of true reports like ‘Mary said that she went to Harvard, but not Harvard’ – where the speaker pronounced the tokens of ‘Harvard’ differently. Here, we agree, that something like mixed quotation might be in play, but there is this difference: the import of the examples need not survive under translation. Suppose we have,

Mary said that she likes tomatoes and not tomatoes.

with different pronunciations indicated. We could not translate this example into, say, German since the relevant pronunciations would be lost. This indicates that the last case is a real instance of mixed quotation (or at least mentioning).

Nothing like this is going on in the ‘nearby’ and ‘left’ examples (or any other of the cases we discuss in IS) - they can all be translated.

b. HR is motivated by Hawthorne's sense that many 'says that' reports have 'the feel of mixed quotation'. That might be, but as a general claim it is irrelevant. Suppose the situation is this: There are clear cases of indirect reports, there are clear cases of mixed quotation, and there are some borderline cases - cases where it's hard to tell which it is. HR would be well motivated just in case we have reason to believe that the relevant cases fall into the latter category. Our sense is that the relevant cases are *not* borderline cases. Consider an example involving 'nearby': Imagine answering a ringing payphone on the street, and a woman's voice says: "There's a river nearby", then she hangs up. Asked what the caller said, you reply: "The woman on the phone said there is a river nearby" It seems perfectly *possible* that the speaker intended to use 'nearby' and not to talk about "nearby" and that in so doing, she succeeded in saying something true (if our earlier diagnosis (of the character of 'nearby') is correct, that's what we would expect). This case, it seems to us, generalizes. The relevant kinds of cases do not *need* to have the 'fell of mixed quotation' (this, again, isn't to deny that it *can* have that feel, i.e. can be uttered with such intentions.)

   This might end up just as a clash of intuitions. There is, however, a rather simple procedure for testing whether your intuitions go with Hawthorne or us: *try doing it*. If there are hidden quotation marks in an indirect report, the speaker must intend for the quotation marks to be there, since the speaker's *intentions* are an important (if not the sole) determinant of that fact. So, the test is simply to try to utter the relevant sentences without quotation marks. If you succeed in this test, you have pretty good evidence that HR is not well motivated.

c. A second point about HR: In the entire literature on mixed quotation, since we introduced the term in (1997b), there's a consensus that mixed quotation is *also* indirect quotation – in particular, the indirect quotation you get by removing the quotation marks. Indeed, many have argued that the sole semantic content of a

mixed quotation is the indirect quotation you get by removing the quotation marks. If this is the correct theory of mixed quotation, the problematic indirect report is still asserted. So the presence of the alleged tacit quotation marks would make no difference.

In the interest of full disclosure, we should mention that we are currently in the search of an alternative to the standard view. As a result, we are well positioned to say that it is extremely hard to find one. Hawthorne, presumably, also does not endorse the standard view, so he owes us a positive theory. Absent such a theory, HR is difficult to evaluate.

Variety of Tests: Agree, suggest, claim, propose, doubt, reject, etc.

Hawthorne asks: 'the authors run the test using the verb 'say'**.** But why are they so preoccupied with this verb? It seems to me to be at least somewhat instructive to try out some other verbs.' He goes on to consider the verb 'agree'. Two quick comments:

a. It's false that we focus on 'said' exclusively in *IS*: throughout we use a variety of terms, such as 'assert', 'claim', 'state' 'ask', and 'order'. We even briefly discuss 'agree' in Chapter 8, but we did not run the tests on them. The reason we did not is that 'say', 'assert', 'claim', 'state' seem to work more or less the same way with respect to the tests. We therefore thought whatever results culled from these would provide an important notion of content. These verbs, after all, are central for our descriptions of the content of others' speech acts. That probably accounts for their prevalence in the literature since Frege. We are unaware of anyone from Frege through Grice through Kaplan up to the present who thinks that 'agree' is a good verb for identifying content.

b. What should we conclude, however, if 'agree' yields different results from 'say', 'assert', 'claim', etc? We're not sure. It's an interesting question, and we plan to pursue it in further research. If 'saying' and 'agreeing' come apart, it will tell us something important about the relationship between attributions of contents and their evaluations. Just what (a form of relativism or something else?) we will suspend judgment on. (See our 'Reply to Critics' (2006) and our 'Shared Content and Semantic Spin' (ms) for further reflections in this issue.)

Conflicting Tests: What to Do?

IS offers another test for context sensitivity – the collectivity test – according to which, if true (simultaneous) utterances of sentences 'a if F' and 'b is F' are collectable into a third context C'' with a true utterance of 'a and b are F', then 'F' is context *in*sensitive. If it can't be truly collected, that's evidence 'F' is context sensitive. Hawthorne says: 'Despite the authors' apparent convictions, their tests do not work in harmony'. In particular, he suggests, our collective description test yields a different result when applied to 'left' and 'nearby' than our indirect disquotational report test does.

Two quick comment: a) We're unconvinced the two tests conflict: for reasons given above, the disquotational report test, as Hawthorne applies it to 'left' and 'nearby', does not instantiate the tests we proposed (both of his own proposals for why we get the context insensitive result implies that the test has been improperly applied). b) If it should turn out that various tests for context sensitivity provide conflicting evidence, we don't see that as a reason for giving up one of the tests. What you do is proceed as in all such cases: you try to develop theories that reconcile the data or you check the evidence again. Think of the diagnosis of the 'nearby' example as an illustration of how to proceed (so, of course, is Hawthorne's mixed quotation suggestion, if you find it convincing.)

'The Appropriate Lesson': Don't Conflate Kaplan's Paradigm Indexicals with the more general phenomenon of Context Dependence

Hawthorne says '…we should all by now have learnt not to conflate Kaplan's paradigm indexicals with the more general phenomenon of context-dependence.' Later he says '[t]he appropriate lesson is this: context-sensitive terms are not alike with regard to the kinds of triggers and signals that allow us to smoothly shift contexts within a single conversation.' W are not committed to the idea that all context sensitive terms are alike with respect to the kinds of triggers and signals that allow context shifts within a conversation. We try, throughout *IS*, to remain neutral about which mechanisms determine the semantic value of a context sensitive expression and the ways in which context determines semantic values. The examples elicited by Hawthorne do not convince us our neutrality was mistaken.

Hawthorne on the 'Awfulness' of Invariantism about 'left'

Hawthorne says:

> Why is semantic invariantism about 'left' so awful? Suppose I am facing you. You say, truly 'The balloon is going left'. I say, truly, 'The balloon is going right' and 'The balloon is not going left'. If we assume there is no context-dependence, then it seems that the following sentences, as uttered by me, should be true:
>
> The balloon is going left and right
>
> The balloon is going left and not going left
>
> The authors do want the semantic content of utterances to connect in interesting and informative ways with our ability to use language. But if sentences such as those above were semantically true, then it is very hard to see how semantic competence is going to provide much of a basis for anything else.

For reasons given above we're not semantic invariantists about 'left', so we will shift to a similar example (one we *are* invariantists about), consider first an utterance u of 'A is red' (where the outside of A somehow is salient in the context of utterance). Suppose u is intuitively true (because A is red on the outside). Consider next an utterance u' of "A is not red" (where the inside of A somehow is salient in the context of utterance). Suppose u' is intuitively false (because A isn't red on the inside). We do not want to infer from this that A is both red and not red, nor that 'red' is context sensitive. Instead, our explanation is that our intuitions about the truth-values of these utterances track the salient speech act content of u and u' –not the semantic content of the sentences uttered. The salient speech act content of u might be *that A is red on the outside*, and the salient speech act content of u' *that A is red on the inside*. These can both be true. The semantic content of u is *that A is red* and the semantic content of u' is *that A is not red*. These can't both be true. Chapter 11 of *IS* is entirely devoted to an explanation of why we don't think it is the semanticist's job to present a theory of what it is to be red, and therefore, why it is not the semanticist's job to tell you which one of u and u' is semantically true. We say these are metaphysical questions, and not semantic, and we give a very long argument for why.

Hawthorne on Semantics, Metaphysics and Ordinary Speakers' Intuitions

Hawthorne asks rhetorically, if 'ordinary people don't know' what it takes for objects to be red,[5] '…is it really very likely that metaphysicians are going to be able to

[5] Hawthorne's example is 'left', since we're not invariantists about 'left' we changed the example - we do this to illustrate the general strategy we use for responding to this kind of object.

settle the matter?' We have no idea how likely it is that *any* metaphysical (or philosophical, for that matter) question will ever be settled. What we are sure about is that if there's an answer to be had, it won't be found by consulting speaker's intuitions about when what's said by utterances of 'A is red' or 'A is red on the outside' are true. Nor will it help to look for a faculty of peculiarly semantic intuitions.[6] That would be awful methodology for many reasons. No one, we hope, thinks you'll discover the conditions under which something is F, by asking people about their intuitions about when 'A is F' is true. Whether the answer to these questions will be found by someone who labels herself a 'metaphysician' or not, we should probably have remained neutral about.

Cappelen and Lepore on Semantic Content

One final, but important remark on this issue: it is not true that we give *no* guidance for how to get at semantic content. Reading Hawthorne's remarks one might get the impression we think semantic content is a mysterious and illusive feature of sentences that we never come into cognitive contact with. That's not our view. All of Chapters 7 and 12 are attempts to show how to focus on semantic content. The guidance we give, however, counsels us not focus on the kinds of examples Hawthorne discusses. Our two primary pieces of advice are these:

a. Think about how to interpret utterances of S under conditions of 'ignorance', i.e. when you know little or nothing about the context of utterance. That will force you to remove 'contextual influence'.
b. Think about collective reports: what utterances of S across many contexts of utterance share; this is reflected in what we call collective reports, i.e. reports of the form 'They all said that p', when 'they' refers to speakers variously situated who have all uttered S.

The examples Hawthorne employs for eliciting intuitions about semantic content are not like this. They are single utterances accompanied with stories about particular contexts of

[6] In this respect, our view is aligned in obvious ways with Williamson's defense of the epistemic theory of vagueness (see Williamson (1994, esp. Chapter 7).

utterance. If we're right, this is not how to access semantic content. The contextual stories block our access to the semantic content.[7,8]

[7] Hawthorne asks whether our view of semantic content prevents 'the semantic content of utterances to connect in interesting and informative ways with our ability to use language'. All of Chapter 12 is devoted to just this very question (in the form that it has been asked by Carston (2002), Recanati (2004) and Stanley and King (2005)). Since Hawthorne doesn't discuss Chapter 12, we're not sure where our disagreement lies.

[8] We hope these remarks also clarify how we accommodate the 'quasi-logical concerns' Hawthorne discusses at the end of his comment. Limitations of space prevent us from pursuing those issues here.