

CONTEXT-SENSITIVITY

Tom Donaldson and Ernie Lepore

[1] INTRODUCTION

We begin with a puzzle. Consider the following two utterances:

(1) I'm Spartacus! [*Said by Spartacus*]

(2) I'm Spartacus! [*Said by Antoninus*]

What Spartacus said was true, and what Antoninus said was not. Yet the two slaves uttered the exact same sentence, so how can this be?

Admittedly, the puzzle is not very hard, and its solution is uncontroversial. The first person pronoun “I” is – to use a technical term – context-sensitive. When Spartacus uses it, it refers to Spartacus; when Antoninus uses it, it refers to Antoninus. So when Spartacus says “I'm Spartacus”, he expresses the true proposition that he, Spartacus, is Spartacus. And when Antoninus says it, he expresses the false proposition that he, Antoninus, is Spartacus. The sentence “I'm Spartacus” expresses different propositions when used by different people.

Another example will help. Contrast these two utterances, made by subjects in a study carried out by experimental epistemologists:

(3) This is a zebra. [*Said by someone while pointing at a zebra*]

(4) This is a zebra. [*Said by someone while pointing at a cleverly decorated mule*]

(3) is true; (4) is false – even though they are tokens of the same sentence. Why? Because the word “this”, like “I”, is context-sensitive. Roughly, an instance of the word “this” refers to whichever object is being indicated, or otherwise

rendered salient, to the audience of the utterance. In (3), the object indicated is a zebra, and so, the token of “this” in (3) refers to that zebra. This animal is the subject of the sentence. So in saying (3), the speaker truly says, of a certain zebra, that it is a zebra. In the case of (4), on the other hand, the indicated object is a mule. In saying (4), then, the speaker says of a certain mule that it is a zebra. Utterance (4) is therefore false.

Enough examples for now – let’s introduce some terminology. “Context”, as the word is used in semantics, is a technical term. For our purposes it suffices to say that the context of an utterance is comprised of all those features of the circumstance in which the utterance is made that are relevant to interpretation. It would be nice to say something more thorough about what contexts are, but this is a hotly disputed topic, so we will do no more than briefly characterise the issues. One debate concerns how structured a context is and what elements it contains: is it just a set of worlds (Stalnaker (1978, 1998))), or something that has many elements, including a set of worlds, set of salient elements, questions under discussion etc. (Lewis (1979), Roberts (1984))? A second debate is whether contexts are subjective mental representations (Hamm, Kamp and Lambalgen (2006)) or objective abstract objects (Steedman and Stone (2006); Devault and Stone (2006, Ch.2)). For our purposes, these issues can be set to one side. A context-sensitive word, phrase or sentence is one which can receive different semantic values (interpretations) when used in different contexts. In particular, a declarative sentence is context sensitive if it can express different propositions when said in different contexts.

We can now begin to form a list of words which are context sensitive in an uncontroversial and straightforward way:

Pronouns: *I, me, my, you, yours, he, him, his...*

Demonstratives: *this, that, those, these...*

Adjectives: *current, present...*

Adverbs: *today, tomorrow, now, here, there ...*

There is little dispute about the context-sensitivity of the words on this list. Most English speakers quickly acknowledge their context sensitivity when it is pointed out to them. It is also usually thought that tense morphemes – like the ending “-ed” which marks the past tense in English – are also context-sensitive (although this has been disputed – see Ludlow (1999)).

Other expressions, however, are the subject of controversy. Indeed, a great deal of dispute revolves around the question of how context-sensitive natural languages are. What’s more, there is disagreement about how context sensitivity should be explained in systematic linguistic theory. We will begin, in sections [2] through [5], by looking at the question, *Which natural language expressions are context sensitive?* We will go on, in sections [6] through [9], by investigating how context-sensitivity should be best accommodated within linguistic theory.

[2] FURTHER APPARENT EXAMPLES OF CONTEXT SENSITIVITY

[2.1] First Example: “On the left”

“On the left” is not as obviously context sensitive as “I” and “this”, but a little reflection suggests forcefully that it is indeed context sensitive. Two New York tour buses pass one another on the road, going in opposite directions past the American Museum of Natural History. On one bus, the guide says, correctly:

(5) The American Museum of Natural History is on the left.

If the guide on the other bus were to say the same thing, all of the passengers would look the wrong way.

“On the left”, it would seem, is context-sensitive. Roughly, in a given context, something is “on the left” just in case its direction from the contextually

salient location is about ninety degrees anti-clockwise from the contextually salient “forward” direction.

[2.2] Second Example: “It’s raining”

John is having a three-way phone conversation with his friends Nitika (in Bombay) and Kurt (in Vienna). The conversation goes like this:

(6) John: What’s the weather like, Nitika?

Nitika: It’s raining.

John: Kurt?

Kurt: It’s not raining.

It seems that Nitika and Kurt could both have uttered the truth – despite the appearance of contradiction. Arguably, when Nitika says “It’s raining,” she says that it is raining *in Bombay*, while when Kurt says “It’s not raining,” he says that it is not raining *in Vienna*. Does it follow that “It’s raining” is context-sensitive?

[2.3] Third Example: “Heavy”

Suppose that Chiara has bought a bike as a present for her niece. She chose this particular bike because, when compared to the alternatives, it was the most lightweight. Having bought the bike, Chiara went to the post office, intending to post the bike to her niece. She was disappointed to find that she would be charged by the kilogram, and the postage cost was enormous. Now consider:

(7) This bike is not heavy. *[Said by Chiara in the shop, admiring its lightweight construction]*

(8) This bike is heavy. *[Said by Chiara in the post office, despairing at the cost of sending the bike by post]*

Arguably, both utterances (7) and (8) are true. In the context of (7), the bike does not satisfy “heavy”, but in the context of (8), it does. A natural thing to say is that – somehow – there has been a shift in the “standard” or “threshold” for heaviness.

[2.4] Fourth Example: “Green leaves”

(9) is clearly context sensitive, because it contains the word “my”:

(9) The leaves on my maple tree are green.

Charles Travis has argued that this sentence is context-sensitive in another, less obvious, way (Travis 1997). Travis asks us to consider Pia, who owns a Japanese Maple tree, with russet leaves. Thinking that the leaves should be green, Pia paints them, and having finished, she says (9) – apparently truly. Shortly afterwards, Pia receives a phone-call from a friend, a botanist looking for green leaves for a study of green-leaf chemistry. Pia offers him the leaves from her Japanese Maple tree. This time, when she utters (9), she says something false. The point is not that “my maple tree” has a different referent on the two occasions – the same tree is referred to both times. Moreover, the tree’s leaves have not changed colour between the first utterance and the second. As Travis likes to put it, it seems that what “counts as” green is different in the two contexts.

[2.5] Fifth Example: Quantifiers

When someone opens his cupboard, peers inside, sighs, and reports:

(10) All the peaches are rotten.

he is unlikely to be reporting on a threat to the world fruit supply: rather, his quantifier “all the peaches” ranges just over those peaches in his cupboard or those in his kitchen, or some similar contextually salient domain. The very same sentence, used in a different context, would have different truth-conditions because the domain of quantification would be different. Arguably, then,

sentences involving quantifiers are context-sensitive, because the domain of quantification is contextually determined (see Westerstahl (1985, 1989) and Stanley/Szabo (2000)).

[3] MINIMALISM, AND MODERATE AND RADICAL CONTEXTUALISM

Multiplying examples of this sort might quickly lead one to conclude that *all* terms in natural language are sensitive to context, and in complicated ways that are hard to describe systematically. Any given declarative sentence, on this view, can express different propositions in different contexts. One might even begin to think it misleading to speak of *the meaning* of a sentence type – on the grounds that the meanings of its various tokens have so little in common. This sort of position is sometimes called ‘radical contextualism’. Travis and Searle are its two chief contemporary proponents, though there are marked tendencies in Carston (2002) and Recanati (2004).

Radical contextualists are at one end of a spectrum of different views. At the other end of the spectrum are the *minimalists*, who deny there is context-sensitivity in natural language beyond what is introduced by the ‘easy’ examples of context-sensitive expressions discussed in section [1]. The minimalist position is defended in Borg (2004) and Cappelen and Lepore (2005). Between the radicals and the minimalists there are a variety of *moderate contextualist* positions.

[4] OBJECTIONS TO RADICAL CONTEXTUALISM

It would be accepted by everyone that (11) is context-sensitive, because it contains the demonstrative “these”:

(11) These leaves are green.

We could say that (11) has no truth conditions *independently of context*. However, most philosophers would say that an adequate semantic theory for English would provide us with a generalisation of the form:

Relative to a context c, (11) is true iff _____

that is, a generalisation which tells us the truth condition of (11) relative to an arbitrary context. One's first guess might be:

Relative to a context *c*, "These leaves are green" is true iff the objects indicated in *c* are green leaves.

This is perhaps a little crude, but it is usually thought that semanticists should seek generalisations of roughly this kind. The characteristic claim of radical contextualists is that this search will fail; motivated by examples like those discussed in section [2], radical contextualists deny that generalisations like this could capture the subtle and multifarious ways in which the truth conditions of English sentences depend on context.

This, then, is the negative claim characteristic of radical contextualism. Problems arise, however, when radical contextualists get around to defining their views positively. In particular, it is difficult for radical contextualists to answer the question, *What does it take for an utterance of (11) to be true?* Here is Travis struggling with this question:

What could make given words "The leaves are green" true, other than the presumed "fact that the leaves are green," is the fact that the leaves *counted* as green on the occasion of that speaking. Since what sometimes counts as green may sometimes not, there may still be something to make other words "The leaves are green" false, namely that on the occasion of *their* speaking, those leaves (at that time) did not count as green.

(Travis, 1997, pp.101-02; see, also, Travis, 1996, p. 457)

This passage suggests an utterance of (11) can be true even when the indicated leaves aren't green, because they might still *count as* green. Apparently, as Travis sees it, the indicated leaves can count as green even if they *aren't*. This should seem strange all by itself, but the key critical point we want to register is that in saying this Travis has actually contradicted the central claim of his radical contextualism. It would seem that on Travis's view, something like the following is true:

(11_{Travis}) Relative to a context *c*, "These leaves are green" is true iff the objects indicated in *c* count as green leaves in *c*.

But this is exactly the sort of generalisation which radical contextualists claim to be impossible!

A second inconsistency charge against radical contextualism arises from Travis's answer to the question: "What, other than the indicated leaves' being green, could ever suffice for the truth of and utterance of (11)?" To see what's gone wrong, note that Travis's answer itself is provided in a particular context. Let's call this 'Travis's Context', or TC for short. According to Travis, the sentences in his paper have their truth conditions determined in TC. In particular, sentences containing the word "green" in Travis's article, must, according to radical contextualism, have the kind of context-sensitive truth conditions radical contextualism claims all sentences have. So the truth conditions for (12):

(12) What could make the given words "The leaves are green" true, other than the presumed "fact that the leaves are green," is the fact that the leaves counted as green on the occasion of that speaking.

are determined (in part) by what counts as green in TC, i.e., the semantic contribution of the expression "green" as it occurs in an utterance of (14) in TC depends on what counts as green in TC. Let's say that anything that counts as green in TC is green_{TC}. But if this is so, (12) then should be read as (12_{TC}):

(12_{TC}) What could make the given words “The leaves are green” true, other than the presumed “fact that the leaves are green,” is the fact that the leaves counted as green_{TC} on the occasion of that speaking.

We hope it is clear this is not what Travis intends to be saying. He doesn't mean to suggest that the leaves counting as green_{TC} is what would make an utterance of “The leaves are green” true in contexts other than TC (and (12) is clearly about contexts other than TC.) The bottom line is that he seems to be trying to *use* “green” as it occurs in (12) in a context insensitive way. Since, according to him, such context-insensitive uses are not possible, it follows that (12) either says something false (i.e., that counting as green_{TC} is what makes utterances of “the leaves are green” true in contexts other than TC) or it employs the word “green” in a way inconsistent with radical contextualism. (The same point applies to our use of “green” on the RHS of (11_{Travis}), and, for that matter, for our attempt to formulate this objection.)

In short, if these arguments are sound, then the radical contextualists have provided no coherent alternative to the standard picture. We turn now to more systematic ways of distinguishing context-sensitive expressions from those which are insensitive to context. In this next section, we will canvass a number of tests for context sensitivity: that is, procedures for identifying context-sensitive expressions.

[5] TESTS FOR CONTEXT-SENSITIVITY

In section [2], we explored a variety of examples which suggest that context-sensitivity in natural language extends well beyond the easy examples we discussed in section [1]. All of these arguments are more or less controversial. Moreover, taken together they push us towards radical contextualism – a position that we have seen to be inherently untenable. What we would like is a *test* with which systematically to distinguish expressions which

are context-sensitive from those which aren't. In this section we will look at some of the tests that have been proposed for this purpose.

[5.1] The Disquotational Indirect Quotation Test (IDR)

Linguistic practices indifferent to context should expose expressions that are not. We will begin by relying on (Cappelen and Lepore 2005).

Disquotation is usually a safe policy in ensuring accuracy in indirect reporting. When Alice says "Bob Dole is Italian," she is usually and most straightforwardly best reported with "Alice said that Bob Dole is Italian." Context sensitive language, however, mandates adjustments. If Bob Dole utters "I'm Italian," the indirect quotation "Bob Dole said that I am Italian" fails to report him (unless the reporter is Bob Dole). Such obvious considerations lead to a test for context-sensitivity: the easier it is to indirectly report an utterance disquotationally, regardless of indifference or ignorance about its context of use, the less likely it is that its constituents are context-sensitive. In short:

IDR: If reporters can easily and truly indirectly disquotationally report an utterance of a sentence S by an agent A, i.e., with 'A said that S,' despite indifference about, or ignorance of, its original context of utterance, then it is unlikely S is context-sensitive.

[5.2] The Collectivity Test (Collectivity)

Collective reports of utterances of a sentence often preserve truth. If John and Herman both utter "China is larger than Russia," we easily collectively report them with "John and Herman said that China is larger than Russia." An exception is when what's uttered is context-sensitive. Suppose we know John and Herman each uttered "Bill left today." We cannot then accurately collectively report them with "John and Herman both said that Bill left today," unless we happen to know that both John and Herman spoke today. The context-sensitive word "today" blocks collection; and in general, context-sensitive expressions

resist collective accurate reporting. Such obvious considerations lead to a test for context-sensitivity based on collecting what others say: The easier it is to collect distinct utterances of a single sentence into a single disquotational report, regardless of indifference about, or ignorance of, their original contexts of use, the less likely it is that the sentence uttered is context-sensitive. In short,

Collectivity: Let u and u' be utterances of S by A and B . If reporters can easily collect u and u' into a single true indirect disquotational report, i.e., with 'A and B both said that S ,' despite indifference about, or ignorance of, the original contexts of utterance, it is unlikely S is context sensitive.

In both tests – and others, cf. (Cappelen and Lepore 2005) – it is presumed that accurate reporting requires reporter and reportees to be expressing the same thought (or proposition) when the latter's utterances easily admit of a disquotational or collective indirect report.

[5.3] Agreement Tests (Agree)

When we apply the IDR and Collectivity tests to expressions like 'immigrate', 'left', 'nearby', 'it's raining', 'local', 'heavy', and 'ready', the tests suggest that these expressions are *not* context-sensitive. For example, an utterance of (13) can often easily be reported with (14), no matter what the environment of the reporter is like, despite indifference about, or ignorance of, the original context of utterance.

(13) It's raining.

(14) John said that it's raining.

Likewise, when John and Herman each utter (13), they can be easily reported collectively with (15) across contexts, despite indifference towards, or ignorance of, John's and Herman's original contexts of utterances.

(15) John and Herman said that it is raining.

'Left' and 'nearby' submit to the same collection patterns. How can an expression behave like this if it is context-sensitive? It has been argued, for example in Cappelen and Hawthorne (2009) and Stanley (2005), that IDR and Collectivity are poor tests for identifying context-sensitive expressions. Space limitations preempt our going very far into this discussion, but we will consider *agreement tests*, which are recommended in Cappelen and Hawthorne (2009) as replacements for the IDR and collectivity tests.

Suppose John and Herman both utter (13) indicating different locations. While they may easily admit of the collective report (15), it is difficult in such circumstances to characterize them with (16)

(16) John and Herman (both) agree that it is raining.

The same applies to some other supposedly context-sensitive expressions, including those containing "nearby", "left", "local", "ready", "enough", etc.. Exploiting such judgments, some authors have recommended introducing Agree-1 and Agree-2 as tests for context sensitivity (in place of indirect reporting tests which they see as flawed).

Agree-1: If A utters S, B utters its negation, and they are not easily reported as disagreeing, say, with 'A and B disagree whether S', then S is semantically context-sensitive (Cappelen and Hawthorne (2009).)

If A says "I'm happy" and B "I'm not happy", no one would report them with "A and B disagree"; this is because "I" is context sensitive.

Agree-2: If A and B both utter S and can be reported as agreeing, say, with 'A and B agree that S', then that is evidence S is semantically invariant across its distinct utterances. If, on the contrary, distinct utterances cannot be so reported, this is evidence S is not semantically invariant across its distinct context of utterance. (Cappelen and Hawthorne (2009), pp. 54-55)

If both A and B utter “John bought a table”, then if they can be correctly reported with “A and B agree that John bought a table,” this is because “table” is context-insensitive.

In short: if speakers in distinct contexts utter S, but can be reported as agreeing, then S is context-insensitive; and if one utters S while the other utters its negation, then S is context-sensitive only if they needn't disagree.

One reason for preferring agreement judgments over indirect reporting ones in ascertaining context-sensitivity is that the former do not admit of distributive readings while the latter may. Agreement and disagreement require co-ordination on a single proposition, while indirect reporting does not.

[6] CONTENT AND CHARACTER

So far, we have been discussing the question, “Which expressions are context-sensitive?” We have yet to discuss how to explain context-sensitivity in systematic linguistic theory. It is to this topic we turn in the next few sections.

We start with the easiest sort of example:

(1) I'm Spartacus! [*Said by Spartacus*]

(2) I'm Spartacus! [*Said by Antoninus*]

As noted back in section [1], the first person pronoun “I” refers to Spartacus when Spartacus uses it, and to Antoninus when he uses it. Here's a possible reaction to the example:

The example shows that distinct tokens of “I” differ in meaning. So the *type* word “I” lacks meaning. It only has a meaning *when used in context*.

This would be an overreaction to the example. There is an important commonality among the semantic interpretations of the different tokens of the

first person singular pronoun: the word “I” always refers to the person using it. Anyone who fails to recognize this feature of the word has not understood it, and so is not a competent user of the word. This suggests that semantic theory should associate a *function* with the word “I”, a function which takes a context as its argument and yields as its value the person speaking in that context. If you like, this function ascribes a meaning to “I” and the referent of each token of the word is the value of the function when evaluated at the relevant context.

Let’s look at our second example:

(3) This is a zebra. [*Said by someone while pointing at a zebra*]

(4) This is a zebra. [*Said by someone while pointing at a cleverly decorated mule*]

We can associate with the demonstrative type word “this” a function, which takes a context as argument and yields the object indicated, or otherwise rendered most salient, in the context. Each token of “this” then refers to the value of the function evaluated at the relevant context.

Now for some terminology: The *content* of a token expression is its semantic interpretation *in its particular context*. The *character* of a type expression is a function which maps each context to the appropriate content. Notice that even non-context-sensitive expressions can be said to have characters – in such cases the character is a constant function. Notice also that whole sentences have characters. The character of the sentence “I was hungry yesterday,” for example, is a function which maps each context *c* to the proposition that the speaker at *c* was hungry on the day before the time of *c*. The terms “content” and “character” come from David Kaplan, who introduced this way of dealing with context sensitivity (Kaplan 1989a, 1989b).

This approach to accommodating context-sensitivity is standard for dealing with all of the uncontroversial cases of context-sensitivity, as discussed in

section [1]. Other (putative) sorts of context sensitivity must be treated differently, as we shall see.

[7] UNARTICULATED CONSTITUENTS

We now return to an example discussed in section [2.2]. Consider:

(17) It's raining. [Said by John Perry's son, in Palo Alto]

Plausibly, the sentence "It's raining" is context-sensitive. In some contexts – as in (17) – it expresses the proposition that it is raining in Palo Alto; in other contexts it expresses the proposition that it is raining in Paris. Not everyone will agree that "It's raining" is context-sensitive in this way, but let's put that question to the side, and instead ask how to explain this sort of context sensitivity, on the assumption that it exists.

Picking up a line of thought from section [6], we might suggest that the source of the context-sensitivity in "It's raining" is that one of the words in the sentence has a non-constant character. "Rain" is the only good candidate, and it seems that the theory will have to go something along these lines:

The word "rain", when used in a context c , refers to the property of raining at the location *contextually salient at c*.

When one uses "It's raining" in a context in which the contextually salient location is Palo Alto, it will express the proposition that it is raining in Palo Alto. On the other hand, when one uses the same sentence in a context in which the salient location is Paris, it will express a different proposition. So far, so good – the theory seems to work well.

One problem for this approach arises when we consider statements like:

(18) It is raining in Palo Alto.

Suppose we use (18) in a context *c* in which the salient location is Paris. Then, according to the proposed theory, the *content* of “rain” in *c* will be the property of raining in Paris. But then, on the proposed theory, the token of (18) would appear to express the proposition that it is raining in Paris, and it is unclear how “in Palo Alto” could play a role in determining the proposition expressed. Something has clearly gone wrong!

So it seems that we need an alternative to the standard content/character approach. When discussing (17), John Perry famously said:

In order to assign a truth-value to my son’s statement [of (17)] ... I needed a place. But no component of his statement stood for a place ... Palo Alto is a constituent of the content of my son’s remark, which no component of his statement designated; it is an unarticulated constituent. (Perry 1986, p. 206)

Here is Perry’s proposal: The underlying relation here is $\text{Rain}(t,l)$ – which holds when it is raining at time *t* in location *l*. (17) specifies one of the arguments for this relation (the “*t*” argument) but no syntactic constituent of (17) refers to a location, so the second argument is left, as it were, blank. The sentence, therefore, expresses an incomplete proposition, a proposition with a hole in it (a “propositional radical”, as Kent Bach puts it in (Bach 1994)):

(19) $\text{Rain}(1^{\text{st}} \text{ September } 1996, _)$

Different tokens of (17) express different propositions because this propositional radical gets “filled in” in different ways in different contexts. On this view, (17) expresses the propositional radical (19), but a particular *token* of (17) could express the proposition:

(20) $\text{Rain}(1^{\text{st}} \text{ September } 1996, \text{ Palo Alto})$

Some part of what is asserted in such a case does not correspond to a syntactic constituent of the sentence used to make the assertion. It is, as one says, “unarticulated” (cf. Perry 1986). It is somehow, in a way never made clear in the literature, “introduced” into the proposition expressed without any direction from the uttered sentence.

[8] BINDING ARGUMENTS AND HIDDEN INDEXICALS

A number of critics have challenged the idea of an unarticulated constituent; in particular, they invoke the so-called binding argument to support the idea that every object in the proposition expressed is put there by an expression in the sentence uttered. We begin with an example. Consider:

(21) Wherever John visits, it’s raining.

On its most natural reading, (21) expresses the proposition that for any location x that John visits, it’s raining at x . On this reading, some sort of quantification is going on. A natural thought here is that where there is quantification there must be variable binding. Accordingly, the adverbial phrase “Whenever John visits” is a quantificational phrase, binding a variable in the part of (21) that follows the comma. On this view, the variable does not correspond to any pronounced component of the sentence, but just the same it must be there. It might seem strange to say that there are syntactic components of sentences which are not pronounced, but in fact this is routine in modern syntax, as a quick look at any introductory textbook will show (see for example Haegeman (1994)).

But what has this got to do with context sensitivity? Hold onto the idea that “rain” is followed by an unpronounced variable, and look again at (17):

(17) It is raining.

According to the view under discussion, consideration of (21) shows that (17) contains an unpronounced variable. But in the case of (17) there is no

quantificational expression that could bind this variable, and so it occurs free. Since it occurs free, its interpretation is fixed contextually. Roughly, in any context of use, the variable will refer to whichever place is most salient in that context of use. On this view, then, the variable in (17) behaves like an indexical (i.e. like “this” or “that”). Because it is not pronounced, it is called a “hidden indexical”. Theories of this kind are called “hidden indexical theories”.

It has been argued, perhaps most notably by (Stanley and Szabo 2000), that unarticulated constituents theories of (17) (as discussed in the last section) cannot be extended to cope with sentences like (21). If this is right, it provides an argument for hidden indexical theories over unarticulated constituent theories. For rejoinders to this “binding argument”, as it is called, see Cappelen and Lepore (2002) and Cappelen and Hawthorne (2007).

It is perhaps worth pausing to draw attention to the key difference between the unarticulated constituents approach and the hidden indexical approach. Both theories agree that (17), uttered in a certain context, will express the proposition that it is raining in Palo Alto, and both theories agree that no *pronounced* component of (17) refers to Palo Alto. The big difference is that the hidden indexical theory says that there is a component of (17) that refers to Palo Alto – an indexical expression which is present in the syntax but not pronounced – while the unarticulated constituents theory *denies* that there is a component of (17) – even a hidden component – which refers to Palo Alto in the context.

[9] THE DYNAMIC LEXICON

It is not uncommon for language users to coin new words “on the fly,” or to invent new meaning for old words. (Clark 1983) gives many examples, including these:

Subjected to the musical equivalent of 72 hours in a dentist’s waiting room, Bradley was apparently in real danger of being the first tourist ever Muzakked to death. (From the *San Francisco Examiner*)

I think that it's across from a quarry. That's the only way I can landmark it.
(Said by a person talking about finding a beach).

I stopped in Perry's for a quick crab. (Herb Caen, meaning a crab that could be eaten quickly)

Examples like these suggest that we should not think that language users just select their words from a fixed lexicon; rather, the lexicon can be modified and extended in the course of conversation. As (Ludlow 2000) puts it, the lexicon is dynamic.

Clark uses this point to explain some putative instances of context-sensitivity. One of his examples is that of denominal adjectives – i.e. adjectives derived from nouns, like 'Hegelian' (from 'Hegel') and 'metallic' (from 'metal'). Many denominal adjectives have well established meanings, but others do not. Clark gives this example:

Churchillian, for example, might mean *with a face like Churchill, smoking a cigar like Churchill, with a speaking style like Churchill*, or any number of other things. In principle, the list is unlimited; in practice, it is limited by what the speaker can assume the addressees know about Churchill, and will be able to see that he is alluding to. (Clark 1992; see also Clark 1983)

If Clark is right about this, then sentences containing the word 'Churchillian' are context-sensitive, but we do not need to account for this by saying that the word is indexical, or by postulating a hidden variable, or by saying that such sentences do not express complete propositions; rather, we should say that the lexicon is dynamic, and the lexical entry for "Churchillian" is different on different occasions.

[10] THREE FURTHER SUPPOSED CASES OF CONTEXT-SENSITIVITY

A summary paper on the literature on context-sensitivity would not be complete without some discussion of contextualism in epistemology, or mention

of contextualist solutions to the liar and sorites paradoxes. However, our treatment of these issues will have to be cursory, because of limitations of space.

[10.1] Contextualism in Epistemology

Well known arguments from epistemology – How do you know you're not being tricked by an evil demon? – might lead an epistemologist to say:

(22) I don't know that I am forty years old.

However, when filling in his census form later that day, the very same epistemologist might say:

(23) I know that I am forty years old.

At this point, he is likely to be accused of inconsistency – an accusation that he can attempt to duck by arguing that “know” is context sensitive. He might say, for example, that there is a shift in which set of epistemic standards are salient between the two contexts.

Contextualism, in epistemology, is the claim that “know” is context sensitive. Contextualists usually employ this doctrine to reconcile sceptical claims (made in the context of an epistemological discussion) with anti-sceptical claims (made in an everyday context). (DeRose 1995) and (Lewis 1996) provide some motivation for epistemological contextualism. For criticism, see (Stanley 2005).

[10.2] Contextualist Solutions to the Sorites Paradox

Shortly before leaving for the airport with my two suitcases, my travelling companion asks, “Are your suitcases heavy?” In fact, the bags differ in weight by only a gram. I reply:

(24) One of my suitcases is heavy; the other isn't.

It seems that this could not be true: my suitcases hardly differ in weight, so how can I truly say that one is heavy and the other not? The example might lead one to accept:

First: If two things differ only slightly in weight, then either neither is heavy, or they both are.

However, it is notorious that claims like *First* lead to problems. Suppose one has a sequence of suitcases. The first weighs 1kg, the last weighs 100kg, and adjacent suitcases in the sequence differ in weight by only 1g. By repeated application of *First*, one can derive a contradiction from the innocuous assumption that the first suitcase is not heavy and the last is, and the assumption that adjacent bags in the sequence differ only slightly in weight.

Contextualists about vagueness, recalling that there are independent grounds for thinking that ‘heavy’ is context-sensitive (see section [2.3]), offer a principle like this in place of *First*:

Second: In a given context, given any two contextually salient objects that differ only slightly in weight, either both objects are in the extension of “heavy” in the context or neither is.

To put it roughly, the line that divides heavy things from the things which are not heavy moves around so that it never lies between two objects that are contextually salient and very similar in weight. The reason why it is improper for me to assert (24) is that both of the suitcases are contextually salient, and they differ only slightly in weight. Views of this kind are defended in (Fara 2000), (Soames 1999), (Raffman 1996), and (Shapiro 2006).

The theory is attractive because it rejects *First* (which leads to contradiction) and also offers an explanation of why *First* is so attractive. Vague predicates appear to be “boundaryless” because the boundary moves in such a

way that it never lies between two contextually salient objects. Wherever you look, there it isn't.

[10.3] Contextualist Solutions to the Liar

Consider sentence (25):

(25) (25) is not true.

What is its semantic status? We have, in answering this question, at least the following three options:

Option One: (25) is true.

Option Two: (25) is not true.

Option Three: (25) fails to express a proposition; it lacks a truth-condition.

Prima facie, options one and two are inconsistent, as are options one and three. It turns out that it is difficult to adopt one of these options *and then stick with it despite further thought*. Whichever option one chooses, one will find oneself inclined to move to a different one.

One might start by reasoning as follows. Suppose for contradiction that (25) expresses a proposition. Then:

“(25) is not true” is true if and only if (25) is not true.

But “(25) is not true” just is (25), so by Leibniz's law we can infer:

(25) is true if and only if (25) is not true.

This is a contradiction – so by *reductio* we can infer that (25) does not express a proposition. This argument gives us a good reason for accepting Option Three.

However, further thought suggests otherwise. We seem to have shown that (25) does not express a proposition. But it would seem to follow that (25) is not true (since only a sentence that expresses a proposition can be true). So we should accept Option Two.

Again, further thought will force us to change our mind. We seem to have given a sound argument for accepting Option Two – i.e. for accepting:

(25) is not true.

But this just is the sentence (25), so we have given a sound argument for (25). It would seem to follow that (25) is true after all, and we conclude that we should accept Option One. Repeating our initial line of argument will take us back to Option Three – and on it goes.

So to repeat, whatever semantic status we ascribe to (25), further thought seems to force us to assign it a different status. The contextualist concludes that (25) must be *changing* status; on this view, (25) is context-sensitive, and as we reason the context changes in such a way that (25) changes status.

The two classic papers advocating contextualist approaches to the liar paradox are (Burge 1979) and (Parsons 1974). Our summary is enough, we hope, to give the reader the barest hint of the motivation for such approaches. For a more satisfying discussion, and further references, the reader is encouraged to look at Beall and Glanzberg (Forthcoming).

[11] SUMMARY

- Some words – including “I”, “this”, “current” and “today” – are uncontroversially context-sensitive.
- It is controversial how far context-sensitivity extends beyond these ‘easy’ examples.

- At one extreme, minimalists argue that the only context-sensitivity in natural language is that introduced by the ‘easy’ examples. At the other extreme, radical contextualists argue that examples of context-sensitivity are so widespread and multifarious that they cannot be accounted for by a systematic semantic theory.
- Radical contextualist positions are hard to state coherently (see [4]), and are at any rate at odds with the results of tests for context-sensitivity (see [5]).
- We have discussed several different ways of explaining context-sensitivity: Kaplan’s character/content approach ([6]); the ‘unarticulated constituents’ approach ([7]); the ‘hidden indexicals’ approach ([8]) and the ‘dynamic lexicon’ approach ([9]). The character/content approach is standard for the ‘easy’ examples of context-sensitivity discussed in [1], while the other approaches are used in more controversial cases.
- It has been argued that postulating context-sensitivity can help us with the liar paradox and the sorites paradox, and that it can be a useful part of a response to scepticism in epistemology.

References

- Bach (1994) 'Conversational Implicature' *Mind and Language* 9 pp. 124-162.
- Beall J. and M. Glanzberg (Forthcoming) 'The Liar Paradox' in *The Stanford Encyclopedia of Philosophy*
- Bezuidenhout, A (2002) 'Radical pragmatics', in Campbell, J. K. *et al.*, (eds.), *Meaning and Truth: Investigations in Philosophical Semantics*. New York: Seven Bridges Press, pp.292-302.
- Borg (2004) *Minimal Semantics*, Oxford University Press: Oxford
- Burge (1979) 'Semantical paradox', in *Journal of Philosophy*, 76:169-198.
- Cappelen, H. and E. Lepore (2002) "Indexicality, Binding, Anaphora and A Priori Truth", *Analysis*.
- _____ (2005) *Insensitive Semantics* Basil Blackwell: Oxford.
- Cappelen, H. and J. Hawthorne (2007) "Locations and Binding," *Analysis*
- _____ (2009) *Relativism and Monadic Truth*, Oxford University Press: Oxford.
- Carston, R. (2002) *Thoughts and Utterances: The pragmatics of explicit communication* (Oxford: Blackwell).
- Clark (1983) 'Making Sense of Nonce Sense', in Clark (1992) and Flores d'Arcais and Jarvella (1983)
- _____ (1992) *Arenas of Language Use* (University of Chicago Press)
- DeRose (1995) 'Solving the sceptical problem', in *Philosophical Review*, 104, 1-52 (1995)

- Devault, D. and M.D. Stone (2006) "Scorekeeping in an Uncertain Language Game,"
- Fara (2000) "Shifting sands: an interest-relative theory of vagueness", *Philosophical Topics*, 28: 45–81. Published under the name 'Graff'.
- Flores d'Arcais and Jarvella (1983) *The Process of Understanding Language* (Wiley, New York)
- Grice (1975) (1975) "Logic and conversation", in *Syntax and Semantics, 3: Speech Acts*, ed. P. Cole & J. Morgan. New York: Academic Press. Reprinted in *Studies in the Way of Words*, ed. H. P. Grice, pp. 22–40. Cambridge, MA: Harvard University Press (1989)
- Haegeman (1994) *Introduction to Government and Binding Theory*, 2nd edition. (Wiley-Blackwell)
- Hamm, Kamp and Lambalgen (2006) "There is no opposition between Formal and Cognitive Semantics."
- Jaszczolt (2000) *The Pragmatics of Propositional Attitude Reports* (Elsevier Science Ltd.)
- Kaplan, David (1989a), 'Demonstratives,' in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (Oxford: Oxford University Press), pp. 481–563.
- (1989b), 'Afterthoughts', in J. Almog, J. Perry, and H. Wettstein (eds.), *Themes from Kaplan* (Oxford: Oxford University Press), pp. 565–614.
- Keefe and Smith (1996) *Vagueness: A Reader* (MIT)
- Lewis, D. (1979) "Scorekeeping in a Language Game"

_____, (1996) 'Elusive knowledge', in *Australasian Journal of Philosophy* 74, 549–567

Ludlow (1999) *Semantics, Tense, and Time: An Essay in the Metaphysics of Natural Language* (MIT)

_____ (2000) "Interpreted Logical Forms, Belief Attribution, and the Dynamic Lexicon", in Jaszczolt (2000)

Parsons C. (1974) 'The Liar paradox', in *Journal of Philosophical Logic*, 3:381–412.

Perry, J. 1986. "Thought without Representation." *Proceedings of the Aristotelian Society Supplementary*. 60: 263-283. Reprinted in Perry (1993).

----- . 1993. *Problem of the Essential Indexical and Other Essays*. Oxford: Oxford University Press.

----- . 1998. "Indexicals, Contexts and Unarticulated Constituents." *Proceedings of the 1995 CSLI-Amsterdam Logic, Language and Computation Conference*. Stanford: CSLI Publications.

Raffman (1996) "Vagueness and context-sensitivity", *Philosophical Studies*, 81: 175-92.

Recanati, F. (2004), *Literal meaning* (Cambridge: Cambridge University Press)

Roberts, C. (2004) "Context in Dynamic Interpretation"

Sainsbury (1996) "Concepts without boundaries", in Keefe and Smith (1996) pp. 251-264

Searle, John, (1978) 'Literal Meaning', *Erkenntnis*, 13, pp. 207-224.

_____ '(1980) The Background of Meaning', in *Speech Act Theory and Pragmatics*, eds., Searle, J., F. Kiefer and M. Bierwisch, Reidel: Dordrecht, pp. 221-232.

_____ (1983) *Intentionality*, Cambridge University Press: Cambridge.

Shapiro (2006) *Vagueness in Context* (OUP)

Soames (1999) *Understanding Truth* (OUP)

Stalnaker, R. (1978) "Assertion"

_____, (1998) "On the Representation of Context"

Stanley, J. (2002) 'Making it Articulated', *Mind and Language* 17 (1&2) pp. 149-168. Reprinted in Stanley (2007)

_____ (2005) *Knowledge and Practical Interests* (OUP)

_____ (2007) *Language in Context: Selected Essays* (OUP)

Stanley, J. and Z. Szabo (2000) 'On Quantifier Domain Restriction', *Mind and Language* 15 (2) pp. 219-61

Steedman, M and M. D. Stone (2006) "Is Semantics Computational?"

Travis, Charles (1985) 'On What is Strictly Speaking True', *The Canadian Journal of Philosophy* vol. 15 no. 2 pp. 187-229. Reprinted in Travis (2008)

_____, *The Uses of Sense*, Oxford University Press: Oxford, 1989.

_____, 'Meaning's Role in Truth', *Mind*, Vol. 100 April 1996, pp. 451-466.

_____, (1997) 'Pragmatics', in Hale and Wright (eds) *A Companion to the Philosophy of Language* (OUP)

_____, (2008) *Occasion Sensitivity: Selected Essays* (OUP)

Westerstahl, Dag (1985) "Determiners and Contexts Sets," in J. van Benthem and A. ter Meulen (eds.), *Quantifiers in Natural Languages* (Dordrecht: Foris), pp. 45-71.

_____, (1989) "Quantifiers in Formal and Natural Languages," in D. Gabbay and F. Guentner (eds.) *Handbook of Philosophical Logic*, iv (Dordrecht: Kluwer), pp. 1-131.