What Are We Doing When We Theorize About Context Sensitivity?

Derek Ball
University of St Andrews

It is as widely agreed as anything in philosophy that some words – "I", "she", "that" – are context-sensitive, and it is as controversial as anything in philosophy whether some other words are – "knows" prominent among them. But what are we agreeing to when we agree that "I" is context-sensitive, and what are we disagreeing about when we disagree whether "knows" is?

We all know an answer to this question: "I" is (and "knows" might be) context sensitive as a matter of its *meaning*; it is *semantically* context-sensitive. But in the absence of further explanation, "semantics" and "meaning" are little more than labels for the problem: what does it mean to say that a word is semantically context sensitive?

There are various kinds of answer to this and related questions in the literature: some rely on some pre-theoretical understanding of meaning – for example, what is said by an utterance of a sentence – claiming that a sentence is context sensitive just in case what is said by utterances of it varies depending on the situation in which the utterances are made; others say linguistics (hence semantics) is a branch of psychology, concerned with the operation of some bit of the human brain, so that semantics must be capturing some sort of psychological fact; according to others, semantics aims at characterizing certain linguistic conventions (and is thus perhaps closer to a branch of sociology); and of course there are other possibilities. This chapter surveys a range of views about the aim of semantics and the nature of semantic fact, with an eye toward connecting some relevant disputes about the nature of context and context-sensitivity. I begin by setting out some presuppositions of the discussion; subsequent sections focus on the various views of semantics and their consequences for our take on context sensitivity.

1 Ground Clearing

In order to keep this paper to a reasonable length, I will focus on one kind of semantic framework in which the titular question can be posed: the "model theoretic" tradition,[1] according to which a semantic theory assigns semantic values – typically set-theoretic entities such as functions – to atomic expressions, and describes composition rules by which the semantic values of

[1] The scare quotes are because it is not clear that models play a very substantial role in natural language semantics; see Glanzberg (2014) for discussion.

complex expressions can be determined on the basis of the semantic values of their component expressions and their syntactic structure. Work on context sensitivity in this tradition takes its cue from the treatment of context in Montague (1974c; 1974a), Lewis (1970), and Kaplan (1977): contexts are ordered tuples consisting of (something like, perhaps *inter alia*) a world, time, and speaker;[2] and semantic values – or at least, one kind of semantic value, which we will call *characters* – are functions from contexts to entities of some other kind – for example, the character of a sentence might be a function from contexts to functions from worlds to truth values. (Since it is possible to be misled by terminology, let me reiterate that I am using "character" and "context" to pick out set-theoretic entities.)

My attention to theories of this kind is mostly for expository convenience; the questions we will consider would arise for most other frameworks that try to treat context-sensitivity. But this sort of framework makes certain questions particularly natural and easy to ask. No one thinks that there is an interesting association in the world, independent of the theoretical activity of linguists and philosophers, between the word "I" (say) and a certain function. Functions are theorists' tools; semanticists are using them to represent something about language.[3] But what are they representing?

We can ask this question about semantic theorizing in general, and also about particular aspects of a semantic theory. The treatment of context-sensitivity that we have just sketched has two main moving parts – the description of an ordered tuple (the context) with elements of certain kinds, and the assignment of characters as the semantic value of expressions – and our key questions will be about the representational role of each of these:

Context Representation What do the elements of context represent? What does it mean when we include a particular parameter (location, say) in the context?

Character Representation What does the assignment of a particular character to an expression represent about that expression?

---

[2] Kaplan (1977) does not explicitly endorse this view of context. In his formal system, he simply stipulates that there is a set of contexts C, but says nothing about the nature of the members of C except that for each c in C, c_A is the agent of c, c_T is the time of C, etc., and that A, T, etc. "may be thought of as functions applying to contexts" (1977: 552), so that (e.g.) c_A is the result of applying A to c. The view that contexts are tuples seems a natural fit with the bulk of Kaplan's commitments, so I will proceed with that view in discussion.

[3] On one view, attributed to Montague by Thomason (1974: 2), languages just are mathematical entities. Perhaps a proponent of this view would claim that the association between "I" and a certain function does exist independently of the activities of theorists; it is, in effect, just a mathematical fact, no more dependent on theoretical activity (and no more in the business of representing something about language use) than the fact that 2+2=4. Though I do not deny that this thought is worth taking seriously, I find it hard to relate it to the practice of semantic theorizing with which I am familiar. I therefore propose to set it aside.

These are the questions that I will sketch some answers to in what follows.[4]

2 The Meaning Perspective

What phenomenon does a semantic theory aim to represent or characterize? We have already noted that there are a variety of possible answers to this question. I want to begin with a relatively straightforward one. Prior to systematic semantic theorizing, we have various semantic concepts: meaning, reference, what is said (by a person or by an utterance), truth, and so forth. The *Meaning Perspective* has it that the job of a semantic theory is to systematize and explain facts about meaning in some pre-theoretical sense – for example, facts about what is said, or about the information communicated by an utterance. For example, Larson and Segal claim that facts about the "actual meanings that [...] expressions have", such as the fact that "The English sentence *Camels have humps* means that camels have humps" are "the primary data that we would want any semantic theory to account for" (1995: 2). And it is clear that Kaplan (e.g., (1977: 492-4)) takes facts about reference and what is said by utterances to be among the facts that semantic theory must explain. (Kaplan introduces the term "content" as a synonym for "what is said" (1977: 500).)

A few clarifications are in order. First, the idea need not be that we are just systematizing our pre-theoretic judgements about some class of semantic phenomena. Plausibly, we should allow that our judgements can be revised in light of theory (for example, in the way judgements about what is said can be revised once one is aware of pragmatic phenomena such as implicature). Moreover, we should want more than a systematization of some data; good semantic theories *explain*. Second, on most conceptions of semantics, the idea is not that we explain *all* of the semantic facts. Rather, at least to a first approximation, we take for granted the semantic facts about atomic expressions, and use them (along with composition rules and syntactic structure) to explain semantic facts about complex expressions such as sentences. (Explaining the semantic facts about atomic expressions is the task of *metasemantics* rather than semantics proper.) Third, we have already mentioned

[4] Considerations of space prevent me from discussing several relevant issues (despite the fact that they are probably essential to understanding context-sensitivity, and are among the areas where the most interesting recent work on issues related to context sensitivity has taken place): notable among them, the extent of context-sensitivity in natural language (most of the discussion will focus on pronouns, and we will generally assume a fairly standard list of context-sensitive expressions), binding and anaphora (our attention will exclusively be directed at so-called *deictic* uses – i.e., uses "whose interpretations are not drawn from the immediate linguistic context" (Nunberg 1993: 12)), and the role of logic in semantic theory (something that played a crucial role in Kaplan's thinking about context-sensitivity, but (as far as I know) very little role in debates about epistemic contextualism).

other approaches to semantics, according to which semantics aims to describe facts about psychology, or about social convention. The proponent of the Meaning Perspective need not deny that facts about meaning in her preferred sense are ultimately psychological (or sociological, etc.) in nature; for example, she might claim that meaning facts are ultimately grounded in (or reducible to, or supervenient on) psychological facts, but still maintain that semantics is its own special science, for nearly all practical purposes independent of psychology (in much the way the study of economics is for all practical purposes independent of physics, even though (at least on a physicalist world view) facts about economics are ultimately a matter of physics).

To fix ideas, let's suppose (with Kaplan) that our theory is designed to make predictions about what is said by utterances of sentences. The observation that motivates Kaplan's theorizing about context sensitivity is that what is said by one and the same sentence varies depending on the situation in which it is uttered. When I utter "I am hungry", what is said is that Derek is hungry; when Jonathan utters the same sentence, what is said is that he is hungry. In order to make predictions about what is said by an utterance of a sentence, we need more information about that utterance; in the case at hand, information about who made it. Let us call a situation in which an utterance might take place – either actual, located in space and time, or possible, the kind of thing that would be located in space and time if it were actual – a *concrete situation*. To a first approximation, contexts represent concrete situations. Following Kaplan (1977: 522-3,546), we should distinguish sentences-in-contexts (or "occurrences") from utterances: utterances are speech acts, events that take place in space and time, while sentences-in-contexts are formal entities – something like an ordered pair of a sentence and a context. The natural (from the Meaning Perspective) idea that we are exploring is that sentences-in-contexts represent possible utterances.

Let's try to fill in the details. We are considering approaches to semantics on which the semantic values of sentences are mathematical objects. One standard assumption is that what is said by a sentence in a context is represented by a function from indices – usually thought of as a possible world, or a tuple consisting of a possible world and other parameters – to truth values; on the assumption that the index is just a world, it will be the function that maps a world to truth just in case what is said by the sentence in the context is true at that world, and to falsity otherwise. So the character of a sentence will be a function from contexts to functions from indices to truth values. Context is an ordered tuple. We need this tuple to give us enough information to generate our representation of what is said. The question then is: what values can we assign to elements of the context to ensure that this is possible? In order to answer this question, we will have to say more about how intensions work.

3 Character and Context from the Meaning Perspective

Our representations of what is said are functions from indices to truth values. We will also assign functions (from indices to other entities) to sub-sentential expressions. All such functions from indices to other entities are known as *intensions*. The result of applying an intension to an index is an *extension*. The extension of a sentence is standardly assumed to be a truth value; the extension of a proper name might be an individual. A fairly standard assumption is that the intension of a sentence, and its extension at a given index, are determined in a systematic way – in the jargon, *compositionally* – by the intensions of its sub-sentential components, their extensions at that index, and the sentence's syntactic structure (though see the discussion of monsters in section 6 below).[5] (These components may include elements that are syntactically realized but unpronounced, and we will use the term "expression" to include these elements.)

We can take the syntactic structure as given. Since an intension, given an index, determines an extension, what we will need from context is enough information to determine the intensions of context-sensitive expressions. And in fact, it is typically assumed that we need less than this; on standard theories, the members of context are *extensions*: a speaker, time, world, and so forth. An intension determines an extension (given an index), but it is not typically the case that an extension determines an intension. (For example, if intensions are functions from worlds to truth values, there are many intensions that map the actual world to Truth.) But context-sensitive expressions are usually held to be a special case. Intensions are needed to make sense of the behavior of expressions in modal contexts, "belief" contexts, and the like. But familiar context-sensitive words are rigid designators: they are, or at least seem to be, unaffected by modal contexts. So it makes sense to give them constant intensions: functions that map every index to the same value. For example, the intension of "I" at a particular context in which I am the speaker might be the function that maps every index to Derek. If we assume that the intensions of context-sensitive expressions are all constant functions, then we can determine intensions on the basis of extensions: if the extension of a context-sensitive expression is e, its intension will be a constant function from indices to e.

With this in mind, there is a simple way to ensure that we have the information we need: we can simply let the members of context be the exten-

[5] I have said that intensions are our representations of what is said, and I am here assuming that those same intensions play a compositional role. But what plays a compositional role need not be identical to the representation what is said, even on the Meaning Perspective; on some views, the representation of what is said can be determined from (but is not identical to) compositional semantic value. See Rabern (2012); Ninan (2010) for discussion. For the sake of ease of discussion, I set this complication aside.

sions of context-sensitive words. (For example, if we are only interested in "I" and "that", the context might be an ordered pair $< a_c, t_c >$, where a_c is a speaker (i.e., the extension of "I") and t_c is an object (the extension of "that").[6]) Call this *the Simple Strategy*.

The Simple Strategy is advocated explicitly by David Lewis (1970: 24, 62-5) and David Braun (1996: 161). And it may seem an attractive view from the Meaning Perspective. After all, the Meaning Perspective has it that our objective is to capture facts about some pre-theoretic notion of meaning such as what is said. In the case of context-insensitive words, it is not part of the job of semantics to explain how atomic expressions get their intensions and extensions. Why should the situation be any different with respect to context sensitive vocabulary? We will return to some possible answers to this question momentarily; first, let's see how the Simple Strategy can be developed.

Character, on the Simple Strategy, is a function that maps contexts to constant intensions (which map every world to a particular parameter of the context); for example, the character of "I" will be a function from contexts to functions from worlds to the first member of the context. Formally, we can write:

(1) $[\![\text{"I"}]\!] = [\lambda c.[\lambda i.a_c]]$
(2) $[\![\text{"I"}]\!]^{c,i} = [\![\text{"I"}]\!](c)(i) = a_c$

(Where $[\lambda c.[\lambda i.a_c]]$ is the function that maps every context c to a constant function from indices to the first member of c.)

Formally, the Simple Strategy is a simple as the name suggests. But what are we representing by describing contexts and assigning characters of this kind? The Meaning Perspective has it that we are representing facts about possible situations in which utterances might take place, and facts about sentences that enable us to make predictions about what would be said by utterances of those sentences in those possible situations. But the Simple Strategist's context will end up being a long and (depending on what context-sensitive expressions there turn out to be) diverse sequence of entities. In what sense does this represent a possible situation?

There is a straightforward way to represent a concrete situation: by giving us enough information to pick it out of the space of possibilities. This is the notion of context familiar from Lewis (1980: pp . 28-9) and others: an ordered triple consisting of a speaker, time, and world: $< a_c, t_c, w_c >$.[7]

[6] Things will need to be more complicated to handle repeated uses of the same context sensitive word, as in "That is not identical to that". Lewis (1970: 62-3) suggests a technical solution (which he attributes to Kaplan); other approaches might see this as a case of context shifting mid-sentence.

[7] See Liao (2012) for doubts about the adequacy of this representation of context.

(Location, time, and world would work just as well for most purposes, but we will assume that the first parameter picks out an agent.) Call this the *straightforward context*. It is clear that context as appealed to in the Simple Strategy is not straightforward context. (If we had only a few context sensitive expressions – just "I" or "here", "now", and "actual" – the Simple Strategy's context would correspond with the straightforward context. But even "that" makes things more complicated.) The Simple Strategy's context represents a speech situation indirectly, by including the semantic values that various context sensitive expression would have in that context. Each member of the sequence corresponds to a particular type of context sensitive expression; for example, the first member might correspond to "I", the second to "that", and so on.

This, then, is the Simple Strategy's answer to the Context Representation question:

Context Representation (Meaning Perspective/Simple Strategy) Context represents a concrete situation by giving the extensions that context sensitive expressions would have if uttered in that situation. Each parameter in the context corresponds to a particular context sensitive expression in the language.

On the Simple Strategy, characters are trivial: they only point us to a particular member (the first, say, or the fourteenth) of the context. This works because we are assigning particular values to the parameters of contexts in a systematic way, depending on facts about the concrete situation that we are trying to represent. There is an interesting question here: why do we assign particular values to the parameters of the context that corresponds to a particular concrete situation? To answer this question would be in effect to give a metasemantic theory for context-sensitive expressions: an interesting project, but not one that the Simple Strategist needs to undertake as long as she can somehow give extensions for the expressions under study in a range of relevant cases.

The Simple Strategy will assign some expressions constant characters – functions that map every context to the same intension – and other expressions variable characters. If an expression is assigned a variable character, that indicates that its contribution to what is said depends on the concrete situation in which it is used. And because intensions are a particular way of representing what is said, and characters are functions from contexts to intensions, characters are defined in terms of our representation of what is said. But characters themselves are doing little work; they serve only as formal devices to retrieve information from context. Simplifying somewhat to put the point bluntly:

Character Representation (Meaning Perspective/Simple Strategy) Characters have no representational significance.

Now I want to turn to two objections to the Simple Strategy. The first objection is empirical: there are data that seem to be a matter of meaning in a pre-theoretical sense that the Simple Strategy just does not explain. The simplest sort of data has to do with infelicity. Suppose I gesture at my favorite chair and say, "He is comfortable", intending to refer to the chair (and not, for example, to make a deferred reference to a person who had been sitting there.) Something has clearly gone wrong, and though one could argue about exactly what it is, it seems clearly to be a matter of meaning of the sort that one might want a semantic theory to capture. Or, to take a subtler sort of example, Nunberg (1993: 34) points out that although "that" and "it" are in some respects very similar, it makes sense to say, upon seeing the face of a certain baseball player, "That's my favorite team," but much less sense to say, "It's my favorite team." Why?

The Simple Strategy has no resources to answer these questions. Note that the objection isn't that the Simple Strategy gets things wrong – one could set up contexts to deliver the right predictions here.[8] Rather, it is that the Simple Strategy is incomplete. There is explanatory work to do that the Simple Strategy is not well equipped to tackle.

The second objection to the Simple Strategy, due originally to Cresswell (1973: 111), trades on the fact that it is an open question exactly which words are context sensitive. If very many (or perhaps even all) expressions are context-sensitive, contexts will become unwieldy; perhaps, if we aim to treat all possible context-sensitivity in natural language, the list will become infinite. So the Simple Strategy threatens to make the task of stating a semantic theory difficult or impossible.[9]

David Lewis (1980) took Cresswell's objection to motivate a shift to the straightforward notion of context, and this clearly requires a corresponding shift in the view of character. Accepting the straightforward notion of context delivers a correspondingly straightforward answer to the context representation question:

Context Representation - Straightforward Context represents a concrete sit-

[8] Perhaps the simplest strategy would be to adopt a representational convention that allowed gappy contexts – contexts that include no extension for certain context-sensitive expressions. See Braun (2005: esp. 621-2 n. 6) for a related idea applied to the use of tuples to represent structured propositions.

[9] Cresswell-style objections may not be fatal to the deployment of the Simple Strategy in limited ways. If we are interested in developing a theory that explains some particular linguistic phenomenon, rather than in capturing the full range of context sensitivity in language, the Simple Strategy may be a useful tool that enables us to abstract away from distracting factors. Neale (2004: 96) calls this style of use of the Simple Strategy *methodological anchoring*.

> uation by giving information that would enable one to pick it out in the space of possible concrete situations.

As Lewis points out, if facts about a concrete situation determine semantic facts, and straightforward context gives us enough information to pick out a concrete situation, then straightforward context should give us enough information to determine semantic facts – we just need to engineer characters that can do the work.

What would such characters represent? One possibility is that they are representing facts about how an expression's contribution to what is said is determined. But this raises a question: more or less every expression contributes something to what is said, and (plausibly) in every case this is determined by facts about the concrete situation. So what is special about "I" and its ilk? Why do we assign proper names (say) constant characters, rather than characters that represent the metasemantic facts about how their contribution to what is said is determined (so that, for example, the character of "Derek" might be something like $[\lambda c.\lambda i.$ the object at the end of the causal chain that leads to a_c's use of "Derek" at t_c in $w_c]$)?[10]

One possible answer is that characters represent another pre-theoretic notion of meaning, one on which different uses of "I" have the same meaning even when they make different contributions to what is said. Although it is plausible that there is such a sense, it is not clear that it is precisely enough delineated to bear serious theoretical weight. (Do all tokens of "that" have the same meaning in this sense? Do they have the same meaning as tokens of "this"? If not, in what does the difference consist? What of unpronounced expressions, such as (on one view) restrictions on quantifier domains?) There seems to be a need for further theoretical work. A natural place to start is Kaplan's thought that character is a "semantical rule" (1977: 520). But there are different ways that this idea might be developed. We turn to these in the next section.

4 The Rules Perspective: Psychology

Rules and rule following are extremely difficult and controversial issues in philosophy, especially in the context of semantics, and we cannot discuss most of the deep questions here. Instead, the aim will be to sketch some conceptions of semantics on which something worth thinking of as a kind of rule enters into the picture, and to discuss some ways these views might relate to controversies about context sensitivity and its representation.

The first type of view has it that semantics aims to capture something about the psychology of language users. On the most prominent version of

[10] The presupposition of the question – that we assign (or should assign) names constant characters – can be questioned; see Recanati (1993: ch. 8).

this style of view (associated with Chomsky and his followers (e.g., Chomsky (1986: ch. 4))) the project is to capture what is represented in or by a certain psychological mechanism – to use Chomsky's technical term, what is *cognized*. In the case of syntax, the traditional view has it that what is cognized is rules, which recursively determine the grammatical sentences of the language. The natural extension to semantics would have it that what is cognized is rules that compositionally determine the semantic facts about sentences (e.g., Larson & Segal (1995: 9-12), Borg (2004: ch. 2)).

There are a variety of ways of developing this kind of view. One strategy is to build on the Meaning Perspective. We often can come to know what is said by utterances in various circumstances. One task would be to characterize the psychological mechanism by which we do this. On the hypothesis that this mechanism works in a broadly computational way, the task would be to characterize the representations that are implicated in the functioning of this mechanism.

Another strategy would be to regard the project of characterizing certain mental representations as supplanting, rather than supplementing, the Meaning Perspective. On this view, pre-theoretical notions of meaning are to be viewed with skepticism, as riddled with unclarity and imprecision of a sort that makes them unsuitable for serious theorizing. *What is cognized* is seen as a more tractable replacement for these notions. Exactly what sort of thing it is that speakers cognize should be treated as an empirical question. One hypothesis, naturally suggested by the psychological focus, is that semantics gives rules associating linguistic expressions with concepts (Jackendoff 2002). Another attitude is that we should remain agnostic on this question until more evidence is in (Yalcin 2014).

What might characters look like on this kind of view? The details may depend to some extent on exactly what we want to represent and how we want to represent it. To keep discussion manageable, I will focus on the view that we aim to represent the mental representations underlying our ability to determine what is said. Suppose we begin by assuming that we are working with the straightforward context. Now we might state the character of "I" in precisely the same way as before. But we are now thinking of this character as doing some representational work: speakers know that when someone uses "I", they are speaking of themselves, and the character of "I" represents this knowledge. This gives an answer to the Character Representation question:

Character Representation - Rules Perspective - Psychology Characters are a theoretical representation of the mental representations that underlie speakers' linguistic competence.

The contrast between the Simple Strategy and the psychological version of the Rules Perspective now under consideration is easier to see with

context-sensitive expressions that do not relate in a simple way to an element of the straightforward context. Suppose that the referent of "that" is determined by speaker intentions. We want to describe what speakers know in virtue of which they can extract a word's contribution to what is said from a concrete situation. So the character of "that" might be something like:

(3) $[\![$"that"$]\!]$= [λc.[λi. the object that a_c intends to refer to with this use of "that"]]

We are now in a position to begin to answer the empirical argument presented at the end of the previous section. Consider the case where I attempt to use "he" to pick out an inanimate object. Now it is a rule, which competent speakers cognize, that "he" cannot be used in this way. Suppose we decide that the consequence of such misuse is that nothing is said. We can represent this knowledge by making the character of "he" a partial function – i.e., a function that maps contexts in which the speaker-intended object is a human male to intensions, and that fails to map contexts in which the speaker-intended object is something else to anything at all. Formally, we can write (using the λ-notation as in Heim & Kratzer (1998: 34-5)):

(4) $[\![$"he"$]\!]$= [λc : the object that a_c intends to refer to with this use of "he" is a human male .[λi. the object that a_c intends to refer to with this use of "he"]]

What of Nunberg's observation that "that" allows deferred reference – i.e., in the case of a demonstrative, reference to something other than the demonstrated object, such as using a demonstrated player to refer to a team – while "it" does not? Nunberg suggests that demonstrated objects play a special role in the semantics of some context-sensitive expressions. For example, "that" picks out an object that is related in some intended way to a demonstrated object; demonstrated objects become "pointers to interpretations" (1993: 38). "It", on the other hand, has no use for a demonstration; as Nunberg points out, "You cannot point at one of the glasses of wine sitting before you at the table and say: 'Now *it's* what I call a good burgundy"' (1993: 34). Instead, "it" picks out an "object that is simply salient in the context or in the consciousness of participants" (1993: 33). The idea that "it" allows deferred reference therefore makes no sense, since there is no demonstrated object for reference to be deferred from.

Now one way to write down a Nunberg-style character for "that" would be:

(5) $[\![$"that"$]\!]$= [λc.[λi. the object that stands in the relation that a_c intends to the object demonstrated by a_c]]

But this downplays the special role of demonstrated objects in Nunberg's system. For Nunberg, interpretation of context-sensitive expressions is a two-

stage process: first, one must identify the features of the concrete situation that are pointers to interpretations, and then one must develop interpretations on the basis of these pointers. A perspicuous representation of what speakers know would separate these two processes, and the obvious way to do this is to make demonstrated objects parameters of the context that can be appealed to in specifying characters. (This would be particularly appropriate if we adopt the Chomskian hypothesis that there is a mental module or faculty dedicated to language, and the identification of demonstrata is an input to this faculty, not performed by the faculty itself.) For example, letting a context c be an ordered quadruple $< a_c, t_c, w_c, o_c >$, where a_c, t_c, and w_c are as before, and o_c is a demonstrated object:

(6) $[\![$"that"$]\!] = [\lambda c.[\lambda i.$ the object that stands in the relation that a_c intends to $o_c]]$

We have now developed context beyond the straightforward context, but not in the direction advocated in the Simple Strategy. Context now provides whatever information the rules we are using characters to represent require. Because the aim is to give a better representation of the rules, call this the *Revealing Strategy*:

Context Representation - Rules Perspective - Revealing Strategy Context represents a concrete situation by specifying those features of it that speakers need to use in applying linguistic rules.

5 The Rules Perspective: Sociology

So far, we have been discussing views of semantics on which semantics aims to characterize facts about the minds of language users. Following Chomsky, some proponents of this style of view draw a sharp contrast with views of language that emphasize social aspects. But a number of theorists have sought to develop the view that language is a social matter – perhaps none more clearly than David Lewis (1983). On Lewis's view, the aim of semantics is to characterize certain social conventions. The semanticist must describe a mapping between sentences and propositions which is such that a speaker s generally make an assertion using a sentence just in case that sentence is mapped to a proposition s believes, and when s hears an assertion made using a sentence, she generally comes to believe the proposition which that sentence is mapped to; and such that these facts are conventional.

It is a difficult question (and one that we cannot consider here) whether the characters of sub-sentential expressions play any representational role in Lewis's system (see Yalcin (2014: 39-42)). But the characters of sentences can play a role in characterizing conventions (though, as Lewis in effect points out, this is not the only way the conventions could be described).

Suppose for the sake of simplicity that an index is just a world. Then the convention might be: if a sentence S has character c, then one must utter S in a concrete situation represented by context u only if one believes the proposition represented by $c(u)$, and if someone utters S in a concrete situation represented by u, one should believe $c(u)$.

On this view, rules are thought of as describing conventions (rather than as describing the contents of certain mental representations). And this gives another possible answer to the Character Representation Question:

Character Representation - Rules Perspective - Sociology Characters represent social conventions.

6 Context-Shifting, Indices, and Contexts of Assessment

Our discussion throughout has tended to assume that contexts are in the business of representing concrete situations, and the discussion has proceeded as though we are thinking of this as a situation in which the expression we are considering might be uttered. But I have been deliberately vague in my "official" statements of answers to the Context Representation Question because there are several factors that may complicate this picture.

First, a number of theorists have claimed that character and context also play a compositional role, because there are what Kaplan called "monsters": operators on character, that shift the context at which expressions in their scope are evaluated. The idea would be that in addition to expressions like "It is necessary that", which have as their extensions functions from intensions to truth values, there are expressions that have as their extensions functions from *characters* to truth values. For example, Kaplan considers (but rejects) the possibility of an operator with the following semantics:

(7) $[\![$"In some contexts it is the case that"$]\!]^{c,i} = [\lambda s : s$ is the character of a sentence. $\exists c'.s(c')(i^{c'}) = 1]$ (where $i^{c'}$ is the index determined by c'; see below for further discussion)

For example, an utterance by me of "In some contexts it is the case that I am hungry" would be true if and only if the agent of some context (not necessarily me) is hungry at the world of that context – i.e., if someone in some world is hungry.

It is debatable whether there are (or could be) monsters in English – Kaplan (1977: 510-2) claims that there could not be (though see Sartorio (2012)) – but a number of theorists have argued that there are monsters in other natural languages (Schlenker (2003)). If there really are operators on character, it is unclear what we should regard contexts as representing. One possible view is that they play a hybrid role: when we begin evaluating an expression, context represents a concrete situation (i.e., that in which we

are considering the expression as being uttered), but as we work through the compositional semantics of the expression, it can come to represent something different (perhaps an aspect of our psychological processing) as it is shifted by operators.

This leaves our answer to the Context Representation question muddled. But as Stalnaker (2014: 214-6) points out, there is an alternative. One can instead insist that context represents a concrete situation. This is in effect to stipulate a representational convention that makes monsters impossible. But one can capture the allegedly monstrous data by building further parameters into the *index*. We can then give the expressions whose interpretations are allegedly shifted by monsters constant characters, functions that map every context to an intension that points to the relevant parameter of the index. For example, we might let indices i be pairs of a world w_i and a speaker s_i, and define a shiftable "I" as follows:

(8) $[\![\text{“I}_{\text{shiftable}}\text{”}]\!]^{c,i} = s_i$

Then we could introduce an operator "In some indices it is the case that" which shifts "$\text{I}_{\text{shiftable}}$" in the way that "In some contexts it is the case that" shifted "I":

(9) $[\![\text{“In some indices it is the case that”}]\!]^{c,i} = [\lambda s : s \text{ is the intension of a sentence. } \exists i'.s(i') = 1]$

This would deliver the result that "In some indices it is the case that $\text{I}_{\text{shiftable}}$ am hungry" is true if and only if someone in some world is hungry, much like "In some contexts it is the case that I am hungry" purports to do. This gives us a mechanism that can make sense of seemingly monstrous data, even while we insist that contexts cannot shift (since the representational role of context is such that context shifting makes little sense).

How does "$\text{I}_{\text{shiftable}}$" behave when it is not in the scope of an index-shifting operator? The extension of "$\text{I}_{\text{shiftable}}$" is fixed by a parameter of the index. So the relevant question is: how do we assign a value to this parameter? Answering this question highlights a role of the context that we have so far ignored. In order to assess the truth of an utterance, we need both a context and an index. And when we evaluate utterances, we typically begin with an index that corresponds closely to the context. For example, on the assumption that an index is just a world, when we evaluate whether an utterance is true, we will begin by letting the index be the world of the context. Suppose that we are evaluating an utterance of "I am hungry now" made by me in the actual world. Then my utterance is true just in case what I have said is true *at the actual world*. This suggest a further role for context:

Initializing the Index Context should give us enough information to determine the index at which we evaluate utterances.

This puts us in a position to consider the second way in which some theorists have questioned the idea that contexts represent concrete situations in which we are considering a sentence as being uttered. Contemporary *relativists* have claimed that we must take into account not only a context in which a sentence might be uttered, but also a context in which it might be assessed. This view can be developed in two ways. First, it might be that some expressions depend on a context of assessment for their intensions. For example, considering the example of a televangelist who says, "Jesus loves you" to an audience widely spread in space and possibly also time (Egan 2009), Cappelen (2009) maintains that what is said depends not only on the concrete situation in which the sentence is uttered, but also on the concrete situation in which the sentence is heard and interpreted (since this determines the contribution of "you"). If this is correct, then we should see contexts as representing at least two distinct concrete situations. (Whether we do this by letting characters be functions from two contexts to intensions, or by maintaining a single context but adding more parameters, looks like a technical matter of little interest as long as we are clear on what we are representing.)

Set Cappelen's *content relativism* aside. The second, and more common, way a relativist view might be developed is as letting features of the index be initialized by the context of assessment. For example, suppose we assume that intensions represent what is said, but let indices be pairs of a world and a speaker $< w_i, s_i >$; and suppose that our representation of what is said by an utterance of 10 is a function from indices to truth values that maps i to truth just in case chili tastes good to s_i in w_i:

(10) Chili is tasty.

The question then is: when we evaluate an utterance of (10) for truth, what index do we use? For example, suppose that I am evaluating Jonathan's assertion of (10), and that Jonathan loves chili and I hate it. Then I can evaluate what Jonathan said at <the actual world, Jonathan>, in which case I should regard it as true, or at <the actual world, Derek>, in which case I should regard it as false. The relativist claims that the latter proposal better fits our actual practice (e.g., because I might regard myself as disagreeing with Jonathan, and might want to argue with him) (MacFarlane 2014). If that is right, then again, we will need to build information representing the concrete situation of assessment into context.

How does this change our proposed answers to the Context Representation question? We could, if we were interested in being fully explicit, state explicitly relativist and non-relativist versions of each of the possibilities described above; but for the purposes of this paper, it is enough if we keep in mind that a complete theory that maintains that contexts are representing concrete situations must describe not only how they do this, but what

situations are represented.

7 Conclusion

One's views about context and character will depend on what we take the task of semantics to be. What is a semantic theory a theory of, and how does it represent its target?

Although some of the questions discussed in this chapter have been debated in some detail (if not in quite these terms) in the literature on epistemic contextualism, others have barely been broached. Do contextualist views give an easy victory to the sceptic (a worry discussed in DeRose (2004))? The answer will depend in part on what we take context to represent. Is there a problem about semantic blindness (as Schiffer (1996) and many others contend)? The answer will depend on what we take characters to represent; semantic blindness manifests quite differently for views that aim to characterize rules represented by speakers than in views that aim to characterize a social convention, or that aim to make predictions about what is said. Attention to the titular question of this paper may not resolve these debates, but it is crucial to understanding what is at stake.

References

Borg, E. (2004). *Minimal Semantics*. Oxford: Oxford University Press.

Braun, D. (1996). Demonstratives and their linguistic meanings. *Noûs*, 30, 145–173.

Braun, D. (2005). Empty names, fictional names, mythical names. *Noûs*, 39, 596–631.

Burgess, A. & Sherman, B., Eds. (2014). *Metasemantics*. New York: Oxford University Press.

Cappelen, H. (2009). The creative interpreter: Content relativism and assertion. *Philosophical Perspectives*.

Chomsky, N. (1986). *Knowledge of Language*. New York: Praeger.

Cresswell, M. (1973). *Logics and Languages*. London: Methuen.

DeRose, K. (2004). Single scoreboard semantics. *Philosophical Studies*, 119, 1–21.

Egan, A. (2009). Billboards, bombs and shotgun weddings. *Synthese*, 166, 251–279.

Glanzberg, M. (2014). Explanation and partiality in semantic theory. In Burgess & Sherman (2014).

Heim, I. & Kratzer, A. (1998). *Semantics in Generative Grammar.* Oxford: Blackwell.

Jackendoff, R. (2002). *Foundations of Language.* New York: Oxford University Press.

Kaplan, D. (1977). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes From Kaplan* (pp. 481–564). New York: Oxford University Press.

Larson, R. & Segal, G. (1995). *Knowledge of Language.* Cambridge, MA: MIT Press.

Lewis, D. (1970). General semantics. *Synthese*, 22, 18–67.

Lewis, D. (1980). Index, context, and content. In *Papers in Philosophical Logic* (pp. 21–44). New York: Cambridge University Press.

Lewis, D. (1983). Languages and language. In *Philosophical Papers Volume I.* New York: Oxford University Press.

Liao, S. (2012). What are centered worlds? *Philosophical Quarterly*, 62, 294–316.

MacFarlane, J. (2014). *Assessment Sensitivity: Relative Truth and its Applications.* Oxford: Oxford University Press.

Montague, R. (1974a). English as a formal language. In Montague (1974b), (pp. 188–221).

Montague, R. (1974b). *Formal Philosophy.* New Haven: Yale University Press.

Montague, R. (1974c). Pragmatics. In Montague (1974b), (pp. 95–118).

Neale, S. (2004). This, that, and the other. In A. Bezuidenhout & M. Reimer (Eds.), *Descriptions and Beyond* (pp. 68–182). Oxford: Oxford University Press.

Ninan, D. (2010). Semantics and the objects of assertion. *Linguistics and Philosophy*, 33, 355–380.

Nunberg, G. (1993). Indexicality and deixis. *Linguistics and Philosophy*, 16, 1–43.

Rabern, B. (2012). Against the identification of assertoric content with compositional value. *Synthese*, 189, 75–96.

Recanati, F. (1993). *Direct Reference: From Language to Thought*. Oxford: Blackwell.

Sartorio, P. (2012). Reference and monstrosity. *The Philosophical Review*, 121, 359–406.

Schiffer, S. (1996). Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*, 96, 317–333.

Schlenker, P. (2003). A plea for monsters. *Linguistics and Philosophy*, 26, 29–120.

Stalnaker, R. C. (2014). *Contexts*. Oxford: Oxford University Press.

Thomason, R. (1974). Introduction. In Montague (1974b), (pp. 1–70).

Yalcin, S. (2014). Semantics and metasemantics in the context of generative grammar. In Burgess & Sherman (2014).