Semantics and Context-Dependence: Towards a Strawsonian Account*

Richard G Heck, Jr

Brown University

Some time in the mid-1990s, I began to encounter variations on the following argument, which was popular with certain of my colleagues and with students who had fallen under their influence.

(i) Semantics is about truth-conditions.

(ii) Only utterances have truth-conditions.

(iii) Hence, a semantic theory must assign truth-conditions to utterances.

(iv) Only a theory that incorporated a complete theory of human rationality could assign truth-conditions to utterances.

(v) So there is no such thing as a semantic theory (in anything like the usual sense).

The argument need not be formulated in terms of truth-conditions. The first premise could equally be stated as: Semantics is about the expression of propositions. The rest of the argument then adapts smoothly. So (i) is meant to be obvious and uncontroversial.

With regard to the second premise, it is important to appreciate the force of the word "only": (ii) asserts that (perhaps with a very few exceptions, such as statements of pure mathematics) sentences *never* have truth-conditions; *only* utterances do. So what underwrites (ii) is not just the existence of context-dependence but, rather, its ubiquity. As has become clear over the last few decades, it is not just the obvious expressions—like "I", "here", "this", and the like—whose meaning seems to vary with the circumstances of utterance. All of the following sentences seem as if they could express different propositions on different occasions, due to context-dependence connected with the italicized expressions:

- *No-one* passed *the test*.
- John is riding *his* bike.
- There are no *tall* gymnasts.

*A shortened version of this paper is to appear in A. Burgess and B. Sherman, eds., *New Essays on Metasemantics* (Oxford: Oxford University Press). Excluded from the published version are sections 3.1, 3.2, 3.3, and 3.5; there are minor changes elsewhere.

Indeed, context-dependence seems to be the norm.

Premises (i) and (ii) are supposed jointly to imply (iii). If sentences do not (ordinarily) have truth-conditions, then semantics cannot be in the business of assigning truth-conditions to sentences. Truth-conditions have instead to be assigned to utterances (or, if you prefer, to sentences relative to contexts). Now, to assign a truth-condition to an utterance, one has to know how the context in which that utterance is made 'resolves' the context-dependence: how the context determines a value for a demonstrative, or a domain for a quantifier, or what have you. As Kaplan would have us put it: We must specify the *character* of each context-dependent expression, where character is a function from contexts to contents. Semantics, then, will issue in statements like:

- "That is a frog" is true with respect to a context C if, and only if, $\phi(C)$ is a frog.
- "Every student passed" is true with respect to a context C if, and only if, every student in the domain $\psi(C)$ passed.

One central task of semantics, on this view, is thus to specify (in an informative way) the functions ϕ, ψ, and so forth that are associated with the various context-dependent expressions.

Premise (iv) asserts that this cannot be done. My early protagonists were particularly impressed by a set of examples concerning color words, such as "This ink is blue". Someone who uttered this sentence would typically mean that the ink would be blue when it had dried on the page, rather than that it was blue when in the bottle, but of course one might also mean that. Or consider "This kettle is black". That might mean that the kettle presently appears black, perhaps because it has been painted; but it might also mean that it would be black (rather than gray) once cleaned of accumulated grime. What proposition a particular utterance of such a sentence expresses seems to depend heavily upon the interests, beliefs, and so forth of the conversational participants. One might conjecture, therefore, that *absolutely anything* could prove relevant to determining which proposition was expressed on a given occasion. In fact, this point had long before been argued in detail by Katz and Fodor (1963). Their topic was ambiguity rather than context-dependence, but their arguments are easily adapted.[1]

In effect, then, what premise (iv) asserts is that no systematic theory is possible of how context affects what proposition is expressed by the utterance of a given sentence.[2] If not, then there are going to be lots of context-dependent expressions whose characters we cannot specify, even in principle. But that, or so said premise (iii), was what we had to be able to do if we were to assign truth-condtions to utterances. So semantic theory is impossible, and that is the conclusion of the argument.

[1]After giving a few examples, Katz and Fodor (1963, p. 178) claim that it is "an easy matter to construct an ambiguous sentence whose resolution requires the representation of practically any item of information about the world. . . ".

[2]Pagin and Pelletier (2007) clearly identify this as the central issue.

This 'anti-semantic' argument has attracted a lot of attention in the last several years, from both friends and foes of semantics. What is perhaps most striking is the way the argument has split the friends of semantics, and the sectarian battles have been bloody indeed. Some philosophers have denied (i), arguing that semantics is interested only in context-independent aspects of meaning (Bach, 2004b). Other philosophers have denied (ii), arguing that, excepting only the most obvious cases, sentences *do* have truth-conditions (Cappelen and Lepore, 2005). The view that seems to be most popular denies (iv), on the ground that the considerations offered in its favor show only that it is *very hard* to say how context affects what is said, not that it is impossible (Stanley, 2007). And, while the friends of semantics squabble, the enemies of semantics point to the squabbling as yet further evidence of the bankruptcy of the entire enterprise. It doesn't help that all three of the options mentioned have serious problems of their own, which defenders of the others are all too happy to point out.

I am going to argue here that this entire debate is unnecessary, because the anti-semantic argument around which it revolves equivocates at a crucial point.[3] There is a sense in which (iii) is true, and there is a sense in which (iii) is false. In the sense in which it is true, it does indeed follow from (i) and (ii); but, in that sense, it does not conflict with (iv); only in the sense in which it is false does it conflict with (iv), but then, in that sense, it does *not* follow from (i) and (ii). So (i), (ii), and (iv) are all true, but they do not entail any anti-semantic conclusion.

Only a fool would try to make this argument in full generality. Fortunately, there is another way to proceed. What I am going to argue is that, if the anti-semantic argument works at all, then it works even if we restrict it to the case of demonstrative utterances, bracketing all other sources of context-dependence. In this restricted case, the question how context affects what is said reduces to the question how context determines the referent of an uttered demonstrative.[4] Premise (iv) then becomes:

(iv′) Only a theory that incorporated a complete theory of human rationality could assign a reference to an uttered demonstrative.

It should be clear that (iv′) contradicts (iii) if (iv) does. So, in that sense, the restriction to demonstratives only makes the anti-semantic argument stronger, if (iv′) can indeed be defended.

One might think it fairly obvious that (iv′) is false. There is a large literature on the question how context fixes the value of an uttered demonstrative; the options are well-known, and they have been explored in depth. In fact, however, I think that (iv′) is not false but true. And if we can establish (iv′), then we will have established, as I said earlier, that the anti-semantic

[3] Of course, there could be other motivations for, say, semantic minimalism, and the arguments here will not address them, if such there be. But, in practice, such views are motivated by the suspicion that we cannot accept all the premises of the anti-semantic argument.

[4] If one thinks, as I sometimes do (Heck, 2002), that uttered demonstratives have sense as well as reference, then the problem that context must solve becomes even harder. So ignoring the Fregean option will not prejudice the discussion.

argument need not have considered fancy cases involving quantifiers, color words, and the like, but could simply have focused on demonstratives. But demonstratives, I shall further argue, do *not* pose a threat to semantic theory, so there has to be something wrong with the anti-semantic argument. What is wrong with it will emerge from our discussion of demonstratives: The same muddle that is responsible for the equivocation at the heart of the anti-semantic argument is present in the literature on how context fixes the reference of an uttered demonstrative; it is, in a sense, responsible for that literature, which I therefore regard as confused. But it will be some time before it will become clear what that might mean.

The plan for the paper is thus as follows. In Section 1, I shall quickly sketch the current state of the literature just mentioned, so as to explain how the dominant ways of approaching demonstratives emerged. The promised discussion of demonstratives is in Section 2, whose central purpose is to argue for a view close in spirit to one Sir Peter Strawson famously expresses in "On Referring": Words do not refer, but people use words to refer. In Section 3, I consider a series of objections to that view. The most important of these will return us to the very general issues we have just been discussing. Indeed, Strawson himself famously believed, at least early in his career,[5] that the 'logical approach to language' that he found in Russell was incompatible with a proper appreciation of the insight just mentioned, and many contemporary philosophers seem to harbor similar suspicions. We shall therefore have to work quite hard to show that the Strawsonian view has little if any bearing upon the prospects of natural language semantics and, in particular, that it poses no threat to the semantics–pragmatics distinction. Finally, in Section 4, I identify a problem left unresolved by the present discussion and suggest an approach to it. Before we begin, though, a warning and a clarification.

First, the warning. We are going to be concerned here with questions about 'what is said' when a given sentence is uttered under certain circumstances. It is important to appreciate that questions about 'what is said' are *not* questions about indirect speech reports.[6] I am not interested in how other speakers would report what someone 'said' by uttering a given sentence. I am not even interested in how they might *correctly* report what had been 'said'. What is at issue here is a question about the content of utterances. There is no particular reason to believe that a sentence of the form "N said that S" can be uttered truly only if N made an utterance whose content, on that occasion, was the same as the content of the sentence S when it is uttered as part of "N said that S". Indeed, there are excellent reasons to think no such strict relationship need obtain. In brief: The notion of what is said is a theoretical notion, one that is intended to play a central role in a general theory of language and its use; there is no reason to think that this notion is expressed by any word of ordinary language, let alone by the English verb "to say". How this notion should properly be characterized

[5]This thought surfaces most strongly in *Introduction to Logical Theory* (Strawson, 1952). Strawson expresses a more conciliatory view in "Meaning and Truth" (Strawson, 1971).

[6]The sorts of points I am making here have been made at length by Cappelen and Lepore (1997). I have made similar points myself about belief and belief-attribution (Heck, 2002, p. 3).

is, of course, an important question, and my remarks here do bear upon it. But, thankfully, it is not essential for our purposes that we answer this question in advance.

Second, the clarification. There is no established use of the word "context" either in semantics or in philosophy of language. I need therefore to clarify how I intend to use the word. So: When I use the word "context", I do not have *any* technical notion in mind. On the contrary, when I speak of 'dependence upon context', what I mean is dependence upon the circumstances in which an utterance is made. This is a wholly intuitive notion.

Many of the technical notions of context that one finds in the literature are representations of context in my sense. They are attempts to extract what we might call the 'relevant features' of the context, so as to make the way in which context determines what is said theoretically tractable. My view is that, at least as concerns demonstratives, but probably as concerns context-dependence generally, this project is hopeless. As I have already said, I strongly suspect that absolutely any fact about the circumstances in which an utterance is made could prove relevant to the determination of what was said. So a context, in this theoretical sense, can be nothing smaller than a (centered) possible world. And there is, or so I shall be suggesting, no prospect of our explaining how the possible world in which an utterance is made fixes its content. That would require a complete theory of human rationality.

There are other technical notions of context one encounters, in which a context is nothing like a representation of the surroundings but is, rather, a representation of the *effects* that context has on interpretation. I have in mind, for example, 'indices', in the sense in which Lewis (1998) uses that term. So far as the present issue is concerned, I have no objection to the use of such a formal tool. If one wants to use such a tool, then the issue I want to discuss is whether it is a critical, or even sensible, question how the circumstances in which an utterance is made (which is what Lewis calls a 'context') fix the indices. I say that the theory of language does not require us to answer any such question. That is not to say that the psychological question how speakers fix the referents of demonstratives might not be worth investigating. Indeed, I am sure it is. It is simply to say that it is a question for psychology, not for semantics.[7]

1 A Brief History of Semantics

Semantics is a child of the revolution in modern logic that Frege instigated. It is therefore no surprise that, for its first several decades, semantics developed without any real attention to the effects of context on meaning. There were, to be sure, squeaky wheels to be heard complaining about the omission. But the initial idealization had to be made; progress would have been impossible without it. And it was not long before serious attention was paid to the more obvious forms of context-dependence. As a result of work

[7] Or for a different part of psychology if, like me, you think semantics is part of linguistics, and linguistics is part of psychology.

by Montague, Lewis, Kaplan, and many, many others, a framework for the incorporation of context-dependence into semantics emerged. Sentences, on this now dominant view, do not express propositions. Rather, they express propositions only relative to a context, which is supposed to supply 'values' for contextual parameters. The values of the parameters then participate in compositional operations, eventually yielding the proposition expressed.

This picture is most natural when applied to so-called 'pure indexicals', such as "I". In this case, it is (at least plausibly) a rule of English that an utterance of "I" refers to its utterer. So, although the language itself does not determine to whom a given utterance of "I" refers, the language and the circumstances in which the utterance is made, taken together, do determine to whom that utterance of "I" refers. There is a sense, then, in which the fact that a particular utterance of "I" refers to me is still an objective fact about my language, even if it is not a purely linguistic fact: Given the facts about the circumstances in which an utterance of "I" was made, the facts about my language determine to whom it refers.

Demonstratives are different, at least in so far as there is no simple rule that determines the reference of an uttered demonstrative. Nonetheless, many philosophers have supposed that there simply *must* be *some* rule that determines to what an uttered demonstrative refers. At least, there must be some such rule if facts about the reference of an uttered demonstrative are to be semantic facts: facts that affect 'what is said' by someone who utters a sentence containing a demonstrative. This attitude derives, it seems to me, from an old and venerable tradition according to which the 'semantic' features of an utterance are those that are determined by the rules of the language. Thus, when attempting to give some explanation of his proprietary notion of 'what is said', Grice (1989, p. 25) writes: "In the sense in which I am using the word *say*, I intend what someone has said to be closely related to the conventional meaning of the words (the sentence) he has uttered". As it happens, Grice is prepared to allow the resolution of context-dependence to affect what is said. Montague (1974) was not: He regarded the semantic features of an expression as including only those that are wholly determined by features of the language to which it belongs. This way of distinguishing semantics from pragmatics is no longer as popular as it once was (King and Stanley, 2004, §1), but its influence nonetheless continues to be felt. Or so I will be arguing.

There are two familiar sorts of answers to the question how the reference of a demonstrative is fixed by the context. Both of these were introduced into the literature by Kaplan, in his discussion of the following now famous example from his paper "Dthat":

> Suppose that without turning and looking I point to the place on my wall which has long been occupied by a picture of Rudolf Carnap and I say: [That] is a picture of one of the greatest philosophers of the twentieth century. But unbeknownst to me, someone has replaced my picture of Carnap with one of Spiro Agnew. I think it would simply be wrong to argue an "ambiguity" in the demonstration, so great that it can be bent to my intended demon-

> stratum. I have said of a picture of Spiro Agnew that it pictures one of the greatest philosophers of the twentieth century. And my speech and demonstration suggest no other natural interpretation. (Kaplan, 1978, p. 239)

So there are two views under discussion. The first, which Kaplan is rejecting, is that the reference of the demonstrative is fixed by the intentions of the speaker; the second, which Kaplan is endorsing, is that it is fixed by what I shall call 'contextual cues', that is, by various aspects of the environment in which the utterance is made, such as: a pointing gesture, or other sort of 'demonstration'; by what is 'salient' in the environment; and so on and so forth.

What I am going to argue is that neither of these views should be accepted, because there is no such problem as the one to which they are offered as a solution. The question was supposed to be: How is the referent of an uttered demonstrative fixed by the circumstances in which the utterance is made? But, to put my view as radically as possible, I do not think there is any such thing as 'the referent of an uttered demonstrative'; uttered demonstratives do not refer at all.[8]

To put the point less radically: I see no reason to believe that the facts about the semantics of English together with the facts about the environment in which a demonstrative expression is uttered do, as Kaplan assumes, conspire to fix a referent for the uttered demonstrative. What I believe instead is very much in the spirit of Strawson (1950, p. 326): that uttered demonstratives do not refer, but that speakers use demonstratives to refer; or again, that, if the reference of an uttered demonstrative is fixed by anything, it is fixed not by the context but by the conversational participants. Better still, what happens when a demonstrative is uttered is that each of the conversational participants will interpret the uttered demonstrative as referring to some object, and will regard himself as having failed to understand the utterance if he is unable to settle upon such an object. Successful communication, in such a scenario, then minimally requires that the conversational participants should converge upon a single object as referent of the demonstrative. If they do, then we may reasonably say that the uttered demonstrative refers to the object on which they converged. But if they do not converge, then, or so I shall be arguing, there will in general be no answer to the question: To what did the uttered demonstrative actually refer? Some of the conversational participants will have taken it to refer to one thing; some, to another. But there is no objective standard that would permit us to say, "This person had it right; that person had it wrong".

It is important to understand the dialectical significance of this claim. The two views Kaplan discusses in "Dthat" can only be distinguished by cases in which they come apart. If there were no such cases, there would be no debate. Cases like the Carnap–Agnew case are supposed to be such cases:

[8]The view for which I shall be arguing is similar in many ways to views for which Kent Bach has been arguing for several years now. See, in particular, his paper "Context *ex Machina*" (Bach, 2004a). But there are also important differences between our views, which shall emerge below.

ones in which the object to which the speaker intends to refer is different from the one the 'contexual cues' plausibly determine as the referent. It takes only a little thought—though I shall return to the point later—to see that judgements about what the 'contextual cues' determine are really judgements about what a suitably attentive person (one who was appropriately aware of the relevant features of the surrounding context) would take the referent to be. If so, then the question which of the two views Kaplan mentions should be preferred just is the question whether the speaker or her audience is correct, and the claim that there is no objective standard that determines who is right implies that neither of Kaplan's two views is right. That makes the problem that Kaplan introduced a pseudo-problem.

As will emerge, and as may be obvious from my continual use of scare quotes, my view is a good deal closer to the speaker's intention view than it is to the 'contextual cues' view. But my view is not a version of the speaker's intention view—not, at least, as that view is normally understood. The view that Kaplan rejects in "Dthat", but later endorses in "Afterthoughts" (Kaplan, 1989a, pp. 582ff), says that it is the speaker's intention that fixes the reference of an uttered demonstrative. Kaplan thus is presupposing that, given a particular sentence of natural language, uttered under particular circumstances, there is some proposition that the sentence expresses relative to the context of utterance; the question is supposed to be what determines which proposition this is; the answer is supposed to be: the speaker's intentions. I mean to be rejecting the question, because I reject its presupposition: Sentences, on the view I want to defend, do not express propositions, *even* relative to the circumstances in which they are made. Rather, an uttered sentence will be *understood* by each of the conversational participants as expressing a particular proposition. The speaker utters the sentence, in part, because she understands it as expressing a certain proposition; the members of her audience will understand it as expressing possibly different propositions; and successful communication is, to first approximation, a matter of everyone's associating the same proposition with the uttered sentence. Which proposition a particular person understands an uttered sentence to express is, in part, a function of how she understands her language—that is the context-independent bit—and, in part, a function of what values she assigns to the contextual parameters. The critical point is again in the spirit of Strawson: It is not context that assigns values to the contextual parameters, but language-users.

I now proceed to argue for this position.

2 Demonstratives

Kaplan's semantics for demonstratives, introduced and developed in his now classic paper by that name (Kaplan, 1989b), incorporates a sort of move that has become quite common. Faced with a completely intractable problem—in this case, what fixes the proposition expressed by an utterance containing a demonstrative—Kaplan separates the problem into a managable bit and an only slightly less intractable bit. The managable bit is found in Kaplan's

theory of *content*: An uttered demonstrative, on Kaplan's account, has as its content an object, and standard compositional operations then yield the proposition expressed. The unmanagable bit is, at least in "Demonstratives", consigned to the theory of character. The character of the demonstrative "that" was to be a function from contexts to contents (that is, objects). And so the problem that remained was to figure out which function that was.

This division does not, however, depend upon the details of Kaplan's own framework. A different approach would take the semantic clause governing the word "that" to entail facts of the following form:[9]

> If u is an utterance of the sentence "that is F", and if the word "that" uttered as part of u refers to x, then u is true iff x satisfies the predicate F.

Very roughly, of course. The question to what the word "that" does refer on any given occasion is simply set aside, in much the way Tarski sets aside the question what sentences mean so as to focus on the notion of truth. Thus, Tarski tells us that an adequate theory of truth for a language $\mathcal{L}$ must imply, for each sentence $S \in \mathcal{L}$, a theorem of the form:

$$S \text{ is true iff } p,$$

where p translates S. But Tarski does not even begin to analyze the notion of translation, and for his purposes he does not need to do so. Similarly, I am suggesting, an account of the *semantics* of demonstratives need not tell us how the semantic value of a demonstrative is fixed. Semantics must tell us how the semantic properties of an utterance are determined by the semantic properties of its significant parts, as well as what sorts of semantic properties the parts have. But to do that, it need not tell us how it is determined exactly which semantic properties the parts do have. That is as true when the parts are context-dependent expressions as it is when they are not. No one would suppose it part of the job of semantics to tell us how it is determined what "snow" means. The important question for semantics is how mass terms function, for example, how the meaning of an expression like "some snow" is determined. There is no more reason, so far as I can see, to suppose that semantics must tell us what determines that a particular utterance of "that" refers to whatever it does. It is enough for semantics to tell us that "that" refers to an object.

But whether or not semantics proper is supposed to answer it, the question Kaplan's semantics for demonstratives leaves unanswered can still seem important: How is the reference of a demonstrative fixed by the circumstances in which it is uttered?[10] The question is particularly important to

[9] King (2001) has argued that complex demonstratives, such as "that dog", are not referential but quantificational expressions, and a similar view might be suggested for simple demonstratives. This question does not matter here, however, since demonstratives are still context-dependent on King's view, so similar questions will arise.

[10] In some sense, it is obvious that semantic theory need not tell us what the referent of an uttered demonstrative is on any given occasion of use. Semantic theory proper knows nothing about the circumstances of use. The most one can hope is that the semantic theory will tell us how to determine the referent, given an appropriate specification of the circumstances. That is Kaplan's hope.

Kaplan, since character is intended as a refinement of the intuitive notion of linguistic meaning, that is, of what a word means in the language to which it belongs. That is to say, the character of "that" is the meaning of the word "that", so we need to specify a function from contexts to contents in order to specify the meaning of the word "that". And, as already noted, Kaplan identifies two options. On the first, the reference of the demonstrative is fixed by contextual cues, the sorts of objective features of the situation in which the utterance is made that someone who heard that utterance might use to assign a value to the uttered demonstrative. On the second, the reference is fixed by the speaker's intention: If the speaker intends to refer to x, then her utterance will refer to x, even if she points at something other than x.

The question which of these views should be preferred has been much discussed, and various sorts of intermediate views have also emerged.[11] But my view, as I have said, is that the question with which Kaplan saddles us is a bad question. It can have no satisfactory answer, because the very asking of it reflects a failure to appreciate something fundamental about how context-dependent expressions work.

Let me first say a few words about the 'contextual cue' view. Part of the difficulty with this view is that no one, so far as I am aware, has ever been able to say, with any precision at all, which features of the environment in which an utterance is made count as 'contextual cues' nor how they conspire to determine reference. Some early writers put a good deal of emphasis upon pointing. I seem to recall someone's even suggesting that we should trace a line from the end of the fingertip and see which object it first intersects. It is obvious upon minimal reflection, however, that pointing is unnecessary for demonstrative reference to succeed. Other sorts of gestures can do just as well—people without arms can refer demonstratively—and in many cases, no gesture at all is necessary. It is therefore common nowadays to think of 'demonstrations' as a means of bringing an object to 'salience', and salience is then taken to be the key to demonstrative reference: An uttered demonstrative is supposed to refer to the object that is most salient when the utterance is made.

One immediate problem is that there are plenty of sentences that contain more than one demonstrative. For example, suppose you and your partner bump into two friends, one of whom is wearing a stunning new necklace. Your partner whispers,[12] "She gave her that for her birthday". This sentence contains three demonstratives. It obviously isn't the case that each demonstrative in the sentence must refer to the same thing. Since there is no upper bound on the number of demonstratives that can be contained in a sentence, we will therefore need something like an infinite sequence of salient objects.[13] How are these to be related to the particular demonstra-

[11]The place to begin, for anyone interested in this debate, are two early papers by Reimer (1991; 1992). The literature grows quickly thereafter.

[12]And does so before you have said anything to them or about them, if 'discourse context' seems to you to be important.

[13]Alternatively, one could suppose that the context changed as each of the demonstratives was uttered. This sort of maneuver is important when dealing with sentences like "It is cold here, but it is warm here", said as one walks from one room to another, but, for reasons that

tives whose references they are? It isn't as if the most salient object will always go with the first demonstrative, and so forth.[14] It really isn't clear what the friends of salience should say here. Still, this is a technical sort of problem, and technical sorts of problems tend to have technical sorts of solutions. But there are more serious difficulties, as we are about to see.

There are two observations about the notion of salience that are worth making right away. First, salience is not an objective notion. What is 'salient' is always salient *to* someone, and what is salient to you may not be salient to me.[15] More importantly, if what is salient determines demonstrative reference, then what is salient cannot be independent of the sentence that is uttered. What is salient will depend upon the meanings of the other words in the uttered sentence, and how it depends upon them will itself depend upon psychological facts about the conversational participants. This can easily be seen by considering the following sort of example.[16]

> Two friends are sitting in the bleachers watching an impressive performance by the local drill corps. Just as one soldier begins the most impressive feat of all, a soldier at the other end of the line makes himself salient by fainting. One of the friends remarks, "He must have been up late".

There is no clear sense in which the soldier who fainted is, *simpliciter*, salient. Suppose, for example, that the sentence uttered had been "He's obviously done a lot of practicing". Then it might have been perfectly obvious to everyone concerned that the demonstrative referred to the impressive performer, despite the fainting. But since the speaker did in fact say "He must have been up late", it is natural to take her to have referred to the soldier who fainted. Which person is salient therefore depends, as was said, upon which sentence is uttered—and not upon its syntactic features but upon its *semantic* features. It is also worth noting, though the point may seem but a curiosity at the moment, that the object to which the uttered demonstrative refers need not become salient until the complete sentence in which the demonstrative is contained has been uttered. In this particular case, for example, had the word "practicing" been followed by, say, the words "to fall on cue like that", then it would again have been the soldier who fainted who was most salient.

We do not even have to change the uttered sentence to get this kind of result. It is easy to add information to our description of the case so that the remark "He must have been up late" would naturally be taken to refer to the performing soldier. We might suppose, for example, that he is well-known to the two friends, that his performance is *way* better than any he has given

will emerge shortly, I doubt that it will help in the present case.

[14] Your partner might have said, "That was her birthday present from her". (There are six options for order.)

[15] This point is also emphasized by Mount (2008, p. 144), whose overall perspective on these issues is close to mine. Unfortunately, I became aware of her paper only after I'd written this one.

[16] This example is apparently due to Kaplan, though I cannot locate a source for it and do not recall where I learned it.

before and, indeed, far better than he had done in the rehearsal yesterday, and that the guy who fainted does so with some regularity. Then what was said would clearly be that the impressive soldier must have been up late practicing to do so well today.

Here is another case, one that has the virtue of being an actual case. Some years ago, I was driving with my friend 'Steve' following a pleasant lunch out. About twenty minutes after we left the restaurant, Steve said, just 'out of the blue', "She was gorgeous". I knew immediately to whom he was referring: the hostess at the restaurant where we'd just had lunch, whom I shall call Sarah, and who looked as if she'd just arrived from a *Vogue* photo shoot.

If salience is the key to demonstrative reference, then we need to understand how Sarah became salient. She was in no way salient to me before Steve's remark—though apparently she was to him. I might never have thought of her again had Steve not said what he did. Nor was she in any way part of the 'discoure context': She had not previously been mentioned by either of us, even when we were at the restaurant. And for that same reason, if one wanted to think of demonstrative reference as involving the selection of an object from some set of candidates for demonstrative reference, then I do not think Sarah was any more a candidate than hundreds of other people were.

Sarah only became salient to me because of what Steve said—and only, again, after he said what he said. And, so far as I can see, there is simply no sense to be made of the idea that she was objectively salient once he had spoken. As I said, I knew immediately to whom Steve was referring, but that is because I'd been at lunch with him, and I too had noticed how strikingly beautiful Sarah was—something he expected me to have noticed. But of course, it is not at all essential that she actually was gorgeous. I would have known just as well to whom he was referring had it just been clear that he was smitten by her. Or she might have been the sort of woman to whom I knew he was particularly likely to be attracted. And that might have been because of her hair color, her dress, the way she carried herself or spoke, or even the sort of car she drove. There are ever so many ways of modifying the example, and it's really just a matter of imagination to bring almost any factor you like into play. So it was only because of what I knew about Steve that I was able to discover to whom he was referring, and it's only because of what he knew about me that he was able to communicate with me as he did.

One lesson to be drawn from this example, then, is that 'contextual cues' can be *anything*, which is to say that there really isn't any sharp notion here at all.[17] Almost any fact about the circumstances in which an utterance is

[17]These reflections lead naturally to the idea, familiar from the work of Stalnaker (1999a), that contexts should be modeled in terms of the 'common ground' of a conversation. In that case, what I am denying is that we need or even can have an account of how the references of demonstratives are determined by the common ground as it is *prior* to an utterance's being understood. (Once the utterance has been understood, it will be part of the common ground, e.g., that Steve was referring to Sarah.) Not that Stalnaker would disagree. His goal is "to redirect attention from the details of the devices and mechanisms of particular natural languages to the general structure of the practices in which natural languages are used to serve the ends of speakers and listeners" (Stalnaker, 1999c, p. 112), and so "to reduce the

made could prove relevant to the determination of demonstrative reference. It is, as I said, just a matter of one's having sufficient imagination to come up with an appropriate sort of example.[18]

One might think this kind of case unusual. But it is actually very common. Some time ago, I was at the aquarium with my daughter when she said, "That's a pretty fish". I looked up and there, high in the corner, was a beautiful blue and yellow fish. I then knew to which fish she'd referred, and I knew it every bit as surely as if she had pointed directly at that fish. It wasn't that it was the most beautiful fish in the tank—though it might have been, and then that would be a different example. In this case, rather, what led me to the right fish was the fact that I knew that this was a fish that Isobel, in particular, would think was pretty. Had it been someone else who'd made the remark—one of a group of tourists, say—I might have had no idea to which creature they'd referred. It therefore seems to me to follow that, if there is a coherent notion of salience to be had at all, it is not one that can be characterized except in terms of what the conversational participants know about one another. If Isobel had hated blue fish but loved orange ones, I'd have understood her differently.

The cue view is motivated by cases like Kaplan's example, cases in which we are inclined to say that the object to which the speaker in fact referred was not the object to which she had intended to refer. As Kaplan puts it, when he points behind him and says "That is a picture of one of the greatest philosophers of the twentieth century", his "speech and demonstration suggest no other natural interpretation" than one on which he said, of the picture of Agnew, that it is a picture of a famous philosopher (Kaplan, 1978, p. 239). But this remark begs the question *to whom* this interpretation is suggested. And the answer must surely be that it is the interpretation that would seem natural to Kaplan's (imagined) audience. That is to say: Our judgements about what the 'cues' fix as the referent are judgements made from the perspective of an imagined audience. But even in this sort of case, our judgements are malleable. Suppose, for example, that it was common knowledge between Kaplan and his audience that the picture of Carnap had been replaced. In that case, he might point behind himself and say, "That picture meant a lot to me", and then he would have said, of the picture of Carnap, that it meant a lot to him.[19] I do this kind of thing myself all the time. I write something on the blackboard and then erase it, only later to realize I still need it. And then, rather than re-write it, I'll say, pointing to where the relevant sentence used to be, "That is the crucial premise in this argument". And my students know perfectly well which sentence I mean.[20]

burden on semantics by explaining as much of the phenomena as possible in terms of truisms about conversation as a rational activity. . . " (Stalnaker, 1999c, p. 113). The present paper is similarly motivated.

[18]As mentioned earlier, this point was made by Katz and Fodor (1963) a long time ago.

[19]Mount (2008, pp. 152ff) discusses some similar cases.

[20]One might try to account for the communicative facts by suggesting that I merely implicate that sentence S is the crucial premise. But my students might argue with me by saying, "No, it isn't", and it is a common observation that one cannot deny mere implicatures that way. For example, if Jones says that there's a gas station around the corner, implicating

One might want to deny that the last two examples really are examples of demonstrative reference, on the ground that they involve what is called 'deferred ostension'. The object to which reference is made is not actually present. But why should that imply that these are not examples of demonstrative reference? Certainly, neither the speaker nor the audience in such a case has what Evans (1982, ch. 6) would have called a 'demonstrative thought' about the referent. So far as language is concerned, however, 'demonstrative reference' is just reference *via* demonstratives, that is, *via* words like "this" and "that", and it is nowhere near obvious that making or understanding a demonstrative *reference* always requires having a demonstrative *thought*.[21] Nor can one simply assume that deferred ostension is somehow exceptional. Indeed, the view that deferred ostension is exceptional seems hard to justify without the prior assumption that ostension is somehow essential to demonstrative reference.[22] But that view has long been discredited, and it should be clear on reflection that both the 'contexual cues' view and the speaker's intention view can happily regard cases involving deferred ostension as genuine cases of demonstrative reference.

If I am right that our judgements about what the 'contextual cues' determine as the referent are, in effect, judgements made from the standpoint of the imagined audience, then it follows that the debate between the cue view and the speaker's intention view is necessarily focused on cases of miscommunication.[23] Any case in which the two views disagree will be one in which the object to which the speaker intends to refer is not the object the imagined audience would 'naturally' identify as the referent. So such a case is a case of miscommunication. But it seems very odd methodology to construct one's theory of how demonstrative reference works by focusing on cases where it fails to work as it should. Surely it would be better to develop our theory of how demonstrative utterances function by looking at what happens when they function properly, and what is required for them to do so.[24]

I will not try to say here exactly what successful communication requires. But I am attracted to the following sort of picture.[25] Speech is a variety of intentional action, which is to say that people generally speak for reasons. That is, people utter the sentences they do to accomplish certain ends, and an utterance's appropriateness to one's end generally depends upon its semantic properties (as well as upon its syntactic and phonological ones). For example, suppose I want to let Chris know that Dave has gone home. One way I might attempt to accomplish this goal is to utter the sentence, "Dave's gone home". But this will make sense only if I assume that Chris will understand me, that is, that he will interpret me as having said that Dave has gone

that it is open, Smith cannot deny that it is open by saying, "No, you're wrong".

[21]This observation undermines some of the examples King (2001, ch. 1) uses to argue that complex demonstratives are not referring expressions.

[22]Mount (2008, §6) argues for a similar conclusion.

[23]Something like this point is also made by Evans (1982, pp. 317–20).

[24]Of course, that is not to say that it might be helpful here to compare successful uses with unsuccessful ones. As, indeed, we will do, to some extent.

[25]I have attempted to develop this sort of view of communication elsewhere (Heck, 2002, 2006, 2007). Related ideas can be found in work by Ian Rumfitt (1995).

home. That is, on my view, Chris will understand me only if he regards me as having spoken the literal truth if, and only if, Dave has gone home.[26] So my reasons for speaking as I do include my expectations about how I will be understood or interpreted. And something similar can be said about Chris's side of the exchange. That he comes to believe (and perhaps even to know) that Dave has gone home clearly depends upon his appreciation of the semantic properties of my utterance, that is, upon his understanding of what I have said: Chris must regard me as having spoken the literal truth if, and only if, Dave has gone home. Successful communication, in such cases, thus depends, minimally, upon agreement between the speaker and her audience about how the utterance is to be interpreted.

Demonstrative utterances are similar. Suppose I see Chris walking towards Dave's office and say, "He's gone home". Here again, I do so with the expectation that Chris will regard me as having spoken the literal truth if, and only if, Dave has gone home. So I expect, *inter alia*, that Chris will regard me as having intended that he would interpret my utterance of the demonstrative "he" as referring to Dave. So, minimally, successful communication that makes use of demonstratives requires that the speaker and her audience converge upon a referent.[27]

If that is to happen, then there is work to be done on both sides. Let's suppose that the speaker wants her audience to come to believe, of some object O, that it is dangerous. Her plan for accomplishing this goal is to utter the sentence "That is dangerous", expecting (a) that her audience will regard her as having spoken truly if, and only if, O is dangerous and (b) that they will suppose she is speaking truly, and so (c) that they will conclude that, indeed, O is dangerous. The speaker is thus presuming that her audience knows a good deal about the words she intends to utter, for example, that the word "dangerous" is true of dangerous things. But, for our purposes, the crucial expectation concerns the demonstrative: She expects that her audience will interpret her as speaking, in uttering it, about O. And she knows, moreover, that her audience will not typically be able to discern this intention telepathically. Rather, her audience will have to rely upon 'contextual cues' to help them identify the object about which she wishes to speak. So it is in the speaker's interest to ensure that her audience will be able to work out to which object she means to refer. Indeed, she cannot

[26] Of course, I must also suppose that Chris will regard me as having spoken truly, but that is of no concern at present.

[27] I keep saying "minimally" not just because there might be other conditions but because I think there are: I think that successful communication requires the speaker and her audience *knowingly* to converge. In particular, I think it needs to be common knowledge between me and Chris that we are both interpreting me as having spoken the literal truth iff Dave has gone home, at least if there is to be a possibility that Chris might come to know that Dave has gone home. But such additional conditions will not matter here.

One might worry, however, that even known convergence is not sufficient. Imagine getting an email saying, "That tree died", and then spending ten minutes trying to figure out which tree the speaker meant. (Thanks to Randall Rose for the example.) Even if one does eventually figure it out, it might not seem as if communication has been very successful in this case. But here we need to distinguish questions. Eventually, the speaker was understood, and so in that sense communication eventually succeeded. But the speaker did not communicate effectively, since there was a long period during which she was not understood.

hope to communicate successfully with her audience unless she ensures that they will be able to identify her intended referent. This, then, is the work the speaker must do: She must ensure that the circumstances in which she makes her utterance will be ones in which her intended audience can determine to which object she means to refer.

Let me repeat that observation, because it is critically important: If one is successfully to communicate using a demonstrative, then one will often have to take steps to ensure that one's audience can determine to which object one means to refer; but, by the same token, one need do no more than is required to make it possible for one's intended audience to discover one's intended referent. It should therefore be no surprise that both how much and what kind of work one needs to do will vary enormously from case to case. A few conclusions are immediate. First, we see why pointing has no special role. Pointing at an object can bring it to the audience's attention and thus lead them to consider it as a possible referent, and that may be enough for them to identify it as the intended referent. But nodding or waving in the general direction of the object can serve the same sort of purpose, as can pointing to where the object was five minutes ago. And no pointing need be required. It isn't just that the object may already be 'salient', though it may well be. The simpler point is that the audience just might not need that kind of help. What help they need will depend upon the circumstances of the case, and there is no limit whatsoever to the sorts of things one can exploit to draw attention to the object about which one wants to speak. That is why absolutely anything can count as a 'contextual cue'.

Note further that the speaker need only make her referential intentions available *to her intended audience*. It is easy to imagine cases where someone might do just that and so purposely conceal her intentions from others who might overhear her words. But in the normal case, it seems to me, one does not purposely conceal one's intentions and yet one will often do too little to allow an arbitrary third party to discern one's intentions. This is precisely what happens in the fish case I mentioned above: I knew which fish Isobel meant, but no one else would have. That is fine, because she was talking to me, and it was of no consequence to her whether anyone else might understand her. The same point applies to the conversation about Sarah, and examples are easily multiplied. It should therefore be clear that, in a very large number of ordinary cases of demonstrative reference, an arbitrary third party would be hopelessly lost. And why not?

The audience has its own work to do. The audience needs to interpret the speaker's words, that is—focusing just on the demonstrative—to fill in the dots in: So-and-so spoke truly iff . . . is dangerous. If the members of the audience are interested in communicating successfully with the speaker, then the question they must ask themselves is therefore: Which object did the speaker expect I would understand her to be speaking about? Or, if you like: To which object did the speaker mean to refer? The question—assuming that it is a different question—is very much not: To what did the uttered demonstrative refer? If the answers to the two questions can come apart, why should you, as a party to the conversation, care whether the 'objective' referent diverges from the one the speaker intended? If you manage to

determine to which object the speaker intended to refer, then the two of you will have converged on a referent, and there is no reason I can see to suppose that you will not then have communicated successfully, even if the object on which you have converged is not the one to which the uttered demonstrative 'really' refers. If so, however, then your failure to converge on the 'objective' referent is of no consequence, and so the 'objective' referent is of no consequence, either.[28]

Undeniably, then, contextual cues have an important role to play in communication *via* demonstratives, but their role is secondary to the speaker's intention. Successful communication requires the speaker and her audience to converge on a referent. But the speaker does not utter the demonstrative and then consult the contexual cues to figure out how to interpret her own words.[29] Rather, in planning her speech, she has already decided what object to assign as value of the contextual parameter that fixes the referent of the demonstrative, that is, which object she intends her audience to interpret her as speaking about. It is up to her audience, if they wish to acquire information from her, to figure out which object that is. Of course, as we have seen, if the speaker wishes to communicate with her audience, then she will need to ensure that they *can* figure out which object she wants to talk about. But the point is that the question the audience must ask is, from the outset, to what the speaker intended to refer. It follows that contextual cues are used as a guide to the speaker's intention.

That is why the notion of salience both seems like it ought to be useful and yet in the end is so useless. When Steve had finished speaking, Sarah was indeed salient. But in what respect was she salient? Only as someone Steve might have wanted to say was gorgeous. So the variety of salience that matters cannot even be characterized except in terms of the intentions with which we speak.

All of this might seem like an argument for the speaker's intention view. I have argued, after all, that the speaker's referential intentions do have a certain sort of primacy. Nonetheless, I do not want to say that the reference of an uttered demonstrative is determined by the speaker's intention. I do not want to state my view that way, because, as I have said, I reject the question to which the speaker's intention view is supposed to be an answer. To repeat, that question is motivated by examples like Kaplan's, in which the object to which the speaker intends to refer differs from the object to which her audience might reasonably take her to refer. The question is then supposed to be to which of these the uttered demonstrative *does* refer. There is no reason to privilege the speaker's point of view over her audience's in such a case. We might try for a synthesis, taking 'what is said' to be determined by some complex calculation based upon a mix of the speaker's intentions and the contextual cues. But where these agree, either will do.

[28]Suppose the uttered sentence was, "That man just stole your wallet". And now suppose that, as the intended audience, you can have the answer to just one of these two questions: To which person did the speaker intend to refer? or: To which person did the demonstrative she uttered refer?

[29]Fodor and Lepore (2004, §4) put great emphasis on this point. We shall discuss their views further below.

And where they diverge, what we have is a speaker who uttered a certain sentence intending to express a certain proposition, and an audience that (reasonably) regards her as having expressed a different proposition. Why exactly do we need to decide what was 'really' said in such cases? I see no point in the question.[30] If so, however, then the question to what some particular token of "that" refers has no point, either, and I see no reason to suppose it has an answer. Not because we do not know or cannot tell, but because it is a bad question.[31]

There is another reason to reject the speaker's intention view, too, one that is less dialectical and more theoretical.[32] It is true that the audience should try to ensure that the referent they assign to an uttered demonstrative is the object to which the speaker intended to refer. But that is not because it is a special fact about *demonstratives* that they always refer to the thing to which the speaker intended to refer. Rather, it is an entirely general principle that, if you want to communicate successfully with the speaker, then you need to ensure that you interpret her words the way she does. This is not only true for other context-dependent expressions, such as color words and quantifiers, but also for words like "livid", about whose meaning speakers regularly disagree (Heck, 2006).

As Strawson (1950) argued, we must distinguish reference as a relation between words and things from reference as an act, that is, distinguish the question to what certain words refer from the question to what a cer-

[30]Brett Sherman suggested I might put this point this way: On my view, intentions are not just one more feature that determines semantic reference; rather, they are constitutive of speaker's reference, which is the only kind of reference there really is. But then one will want to know whether my view has problems dealing with the phenomena for which the distinction between speaker's reference and semantic reference was introduced, and we shall only address those questions below, in §3.6. Still, Sherman's way of stating the view does highlight some of its important aspects.

[31]In "Domain of Discourse", Gauker raises the question how, when speakers disagree about what the domain of discourse is, they might resolve their (essentially verbal) disagreement. Gauker concludes his discussion of what he calls "neutral" solutions—of which mine would be an example—with these remarks:

> My concern is not that we cannot explain in any way the course of events that results in Tommy and Suzy's ceasing to argue. No doubt we would be able to explain it, and would even be able to explain it *in terms of each party's thoughts and intentions*. What we cannot do if we adopt a neutral solution is explain it as a process in which a sequence of utterances each expresses a unique proposition. (Gauker, 1997, p. 29, my emphasis)

I agree, but I do not see why this is supposed to be a problem. Gauker had earlier remarked:

> It is important to bear in mind that our aim in theorizing about the determinants of domains of discourse is not merely to explain linguistic behavior. In order to explain linguistic behavior perhaps we never need to know what the domain of discourse governing a conversation really is but only what each interlocutor *takes* it to be. (Gauker, 1997, p. 27, his emphasis)

Indeed. What is the other aim, though? Gauker does not say, but it would appear that he is supposing that the aim is determining what "unique proposition" is expressed by each utterance. On my view, however, utterances do not express unique propositions; they only express propositions to or for speakers.

[32]Special thanks to Seth Yalcin here.

tain person refers.[33] And once we do, then my view is akin to Strawson's: Words—in this case, uttered demonstratives—never refer; rather, speakers use demonstratives to refer. But one might well worry that, in rejecting reference as a relation between words and things, we thereby reject semantics. Strawson's own skepticism about the applicability of formal methods to natural language might reinforce that concern. We'll return to that issue after discussing some others.

3 Some Objections

There are several reasons one might want to resist the conception of demonstrative reference that I have elaborated. These fall into several broad categories. There are worries about the lack of any role for an objective notion of reference in my account (§3.1). There are concerns about what the account implies about the relation between language and thought (§3.2). There are worries about whether I have over-analyzed what is involved in demonstrative reference (§3.3). There are questions about what this view implies about the prospects of natural language semantics (§3.4) and, in particular, about the principle of compositionality (§3.5). And there are objections that derive from views about the distinction between semantics and pragmatics (§3.6). But before I address any of these objections, let me first dispose of two that seem to me simply to register misunderstandings.

An objection some people may have been wanting to lodge for a long time is that I am confusing the metaphysical question what objectively determines the reference of a demonstrative with the epistemological question how people go about determining the reference of a demonstrative. But I am not confusing these questions. On the contrary, I am consciously running them together. My view is that there isn't any such thing as "what objectively determines the reference of a demonstrative". There are language-users who make and interpret demonstrative utterances, and I am sure that there are interesting things to be said about how they do so, in particular, about how they converge on a referent and so how they manage to communicate successfully. These psychological and epistemological questions are, in my view, the only ones there are here. Once they have been answered, there just isn't another, metaphysical question that remains unanswered: the one about what the referent of the uttered demonstrative really is and what determines it.[34] Such a question is unmotivated, of no obvious interest, and has no clear content.

Another blunt objection is that, if someone right now just says, out of the blue, "That is so beautiful", then no matter how fervently he may intend to refer to Picasso's "Child with a Dove", he can't. Since all views that privilege the speaker's intentions will imply that he can, then there is something wrong with all such views, including mine. This I flatly deny. Suppose that

[33] Of course, the distinction would be pointless if referring (as an act) was making it the case that one's words referred. The worry that it must be will be discussed below.

[34] To be clear, I do not mean to deny that the question, say, that makes it the case that the speaker's referential intention concerns a certain object is not worth asking or can have no answer. The question I am rejecting is specific to language.

Fred has been thinking to himself for some time about "Child with a Dove", studying the image he has before his mind's eye. Lost in contemplation, he says, "That is so beautiful". Then I think he plainly does refer to "Child with a Dove", and I see no point in our saying otherwise. Other people may have a hard time knowing what he means, but that is a different matter. And there is nothing abnormal about this case. People do this kind of thing all the time; soliloquy is a perfectly acceptable form of speech. And if what the objector wants to deny is that Fred's *words* refer to "Child with a Dove", then I am happy to agree. I do not think Fred's words refer at all.

3.1 Objectivity

I said a moment ago that I regard the question what determines the referent of a demonstrative as being of no interest. But someone might want to insist that the question is not just interesting but important. We need the notion of what an uttered demonstrative really refers to, this objector wants to say, in order to explain our ability to communicate with one another. The thought here is structurally similar to one that motivates the view that public languages, such as English, should be our primary focus when we are theorizing about language, rather than idiolects. On the idiolectical view, each speaker has her own language, and successful communication depends upon there being an appropriate sort of overlap between different speakers' languages; more precisely, it depends upon speakers' agreeing about what various expressions mean. But many philosophers have worried that this sort of view will make it impossible to explain our ability to communicate with one another.[35]

Suppose, for example, that I am walking down the street and ask someone the time. When I do, I am supposing that she will understand my question. And when she replies, "It is about noon", and I come to believe that it is about noon, I am supposing that her words mean to her what they mean to me, namely, that it is about noon. What possible reasons could I have for such suppositions? None, the thought goes, unless I have reason to suppose that my interlocutor speaks English, just as I do. That is not, of course, to say that her idiolect must actually be English. Both her idiolect and mine may diverge from English in various ways, since each of us may and presumably does misunderstand various words. But what explains our ability to communicate is the fact that we are both trying to speak English, which is a public language that acts as a kind objective standard for us. That, indeed, is what it is to be a speaker of English: to hold one's own, sometimes idiosyncratic, use of language accountable to the objective standard that is established by the English language. And, this line of thought then concludes, it is only because there is such a standard, only because we are all aiming at the same target, and so all manage to get somewhere in its vicinity, that we can communicate the way we do.

[35] My favorite exposition of this line of thought is due to Wiggins (1997). Such considerations surface frequently in Dummett's work, too. For references, see my paper "Idiolects" (Heck, 2006), which defends the idiolectical view against these sorts of criticisms.

A similar line of thought might seem to be equally compelling in the present case. People are very often able to communicate using demonstratives, which is to say that they are able to converge on a referent. If that is not to be a miracle, doesn't there need to be some objective fact about what the referent is, a fact that we are all trying to get right? And if we are actually to explain our ability to get it right, don't there need to be rules that determine what the referent is, rules that we all in some sense know and then apply?

One reason to reject this demand for an objective standard is that it would prove far too much. We do not just manage to communicate successfully with one another when we speak literally. We also enjoy widespread communicative success when we use non-literal forms of speech, such as implicature and even metaphor. So Professor Smith writes, "Mr. Jones is punctual and has good penmanship". Does it have to be an objective matter of fact that his writing what he did implicates that Jones is no good at philosophy if we are to explain how he can communicate what he does? Do we all have to know some fixed set of rules by means of which we calculate implicatures? I do not see why (Heck, 2005, §1). It is enough that we should be able to ask ourselves such questions as why Smith wrote such a strange letter (that is, with what intentions he did so), that we should pretty good at figuring that sort of thing out, and that we should be good enough at doing so that Smith can rely upon our being able to do so in forming the intentions with which he speaks. And we *are* pretty good at discerning the intentions with which people act. Our existence as social animals depends upon it.

The case of metaphor is even stronger. If Smith writes, "Mr. Jones's mind is a collection of souveniers gathered during ventures into the wilderness", does it have to be an objective matter of fact that what he wrote means that Jones is too disorganized to succeed in graduate school if we are to explain how he can communicate what he does? Do we all have to know some fixed set of rules by means of which we calculate the significance of metaphors? I see no reason to suppose we do.

And for much the same reason, I just don't see any reason to believe that we all have to know rules that determine what the reference of a demonstrative is; nor that there have to be any such rules; nor that there has to be an objective fact about what the reference of a demonstrative is that we are all trying to get right. In the case of the general issue about idiolects, I can at least feel the pull of this sort of worry, though, in the end, I think the pull should be resisted (Heck, 2006). But I do not at all feel the pull of the corresponding worry here. As I emphasized above, a speaker who utters a demonstrative knows that he will not be able to communicate successfully with his intended audience unless he makes it plain to them to which object he means to refer. It is therefore in his interest to do whatever he needs to do to make it possible for his audience to discover his intention. Since people very often do want to communicate with one another, they very often do take steps to make their intention plain. And since we are generally pretty good at discerning each other's intentions, it is no surprise that we are able to do so in these sorts of cases, too.

Indeed, as I mentioned earlier, our interest, in cases of communication, is

ultimately in which object the speaker meant to be talking about. This point is perhaps even clearer if one thinks about questions. If someone asks, "What is in that glass?" wondering, say, if it is safe to drink, then, if I am going to answer their question in a way that will lead them to form an appropriate belief, what I need to know is which glass they mean to be asking about. If, as my opponent believes, there is a possibility that the words "that glass" might actually have referred to something other than what the questioner intended, then it is unclear why I should care about this 'objective' referent. If the speaker failed to secure the right glass as the 'objective' referent of her utterance, but it were clear to me which glass that was, then I would answer her just as if she had secured it as the 'objective' referent.[36] The alleged failure is communicatively irrelevant. But then it is unclear what work the 'objective' referent is doing.[37] And the fact that communicative success only depends upon the speaker's making it possible for her intended audience to determine to which object she means to refer seems to me to make it likely that speakers very often will fail to secure the right object as the 'objective' referent. Of course, it is hard to know, since, as mentioned some time ago, we know neither what this 'objective' referent is supposed to be nor how it is supposed to be determined. But it is clear enough that making one's referential intentions clear to one's intended audience often requires one to do much less than would be required to make them clear to an arbitrary third party, and that suggests to me that one often will not do what is required to secure an 'objective' referent for one's words, whatever that might be supposed to mean.

I am insisting, then, that the facts about communication can be explained in terms of the speaker's intention to refer to a certain object and her audience's attempts to discover that intention. No appeal to an 'objective' referent is required.

3.2 The Autonomy of Language

There is another concern that is in some ways similar to the ones we have just been discussing but whose source is very different. This concern is rooted in a commitment to the priority of language over thought. One might reason as follows. If language is prior to thought, then we must account for the meanings of utterances without making any use of mentalistic notions like belief and intention. Hence, the signficance of an utterance of "That is a frog" simply *must* be a function of objective facts about the language and non-mental features of the context. Since my view denies this, it is committed to the priority of thought over language, so there must be something wrong somewhere.

Now, frankly, I do not myself regard this as an objection. I am quite comfortable with the idea that animals and pre-linguistic infants have thoughts, and in so far as Augustine suggests that language acquisition involves chil-

[36] Suppose I ask first, to make sure I know which glass is in question. Is it possible for the speaker to answer my question wrongly?

[37] There is a parallel argument that can be made in the case of idiolects (Heck, 2006, pp. 73ff).

dren's learning labels for concepts they already possess,[38] I suspect he was to some extent correct. But this view, though I do hold it, seems to me to be independent of the position for which I have been arguing here. That is: I think it would be wrong to object to my treatment of demonstrative reference on the ground that it entails the priority of thought over language, because it just *doesn't* entail it, in any substantial sense.

It should be pretty clear that the position defended here is compatible with views according to which neither language nor thought is prior to the other. Such views have been held by Davidson and Dummett, among many others. On these views, the ability to think about objects is supposed to be co-eval, somehow, with the ability to refer to objects. I don't for a moment claim fully to understand what that means or how it is supposed to work, but Davidson is quite explicit both (a) that thought without language is impossible (Davidson, 1984c) and (b) that our account of the significance of language must simultaneously be an account of the significance of thought (Davidson, 1984b). Meaning, belief, and intention all have to be treated together. The idea that the ability to refer demonstratively to objects, or to interpret others as doing so, is inextricably entangled with the having and discovering of certain sorts of intentions thus seems to fit quite happily with Davidson's views.

That is not to say that the account of demonstrative reference I am defending is compatible with every possible view about the relation between language and thought. I am assuming that we can distinguish the ability to *think of* an object demonstratively, on a particular occasion, from the ability to *refer to* that object demonstratively, on that occasion; more precisely, the former ability cannot be constituted by the latter, on that particular occasion.[39] That makes my view incompatible with views that hold that thought is possible only through language. Gauker, for example, denies that speakers' intentions can play any role whatsoever in resolving context-dependence. This denial issues from Gauker's rejection of the 'received view' that, as he puts it in *Words Without Meaning*, "the central function of language is to enable a speaker to reveal his or her thoughts to a hearer" (Gauker, 2003, p. 3) or, as he put it in an earlier paper, that "linguistic communication is basically a matter of a speaker's choosing words that will convey the propositional content of [one's] thought to hearers" (Gauker, 1997, p. 2). But it does not trouble me to be committed to that view.[40]

[38] As Wittgenstein famously suggests at the opening of the *Philosophical Investigations* (Wittgenstein, 1968).

[39] The point of this qualification is to allow, as above, that, quite generally, the ability to think of objects might depend upon the ability to refer to them. It is just that this cannot hold case by case.

[40] That said, as Rescorla has pointed out, plenty of people otherwise quite unsympathetic to Gauker's position would want to insist "that the primary function of assertion is not to reveal anything about one's own mental states, but rather to describe the subject matter of one's assertion, which typically will be both extra-linguistic and extra-mental" (Rescorla, 2006, p. 121). I would be one of those people.

3.3 Children

A related but distinct worry is that my account overly intellectualizes demonstrative reference in a way that makes communication with children a mystery. Of course, the focus here has been on communication among linguistically competent, relatively high-functioning adults, but the objection is nonetheless in order.

Children who are just learning to speak, the thought would be, are not capable of directing attention, or of estimating intention, in the ways I have suggested are typical of adult communication involving demonstratives. Perhaps so.[41] In that case, the natural hypothesis would be that children have some relatively simple method of determining demonstrative reference. We do not know what that is, so let's suppose that young children always interpret an uttered demonstrative (theirs or another's) to refer to the object at the center of their attention.[42] If so, then that will affect how one communicates with a child. The question one must ask oneself, when a young child speaks, will be not to what he intended to refer, but just what he was thinking about; when one speaks to the child, one will need to bring one's intended referent to the center of the child's attention, since she will always interpret it as referring to what is at the center of her attention.

But none of that seems at all implausible. One *does* have to work harder to communicate with children who are just learning to speak, since one has to do most of the work oneself. And the semantics of demonstratives can still be the same for children as for adults: A demonstrative refers to an object.

3.4 Semantics

We can now return to the anti-semantic argument with which I began this paper:

(i) Semantics is about truth-conditions.

(ii) Only utterances have truth-conditions.

(iii) So a semantic theory must assign truth-conditions to utterances.

(iv) Only a theory that incorporated a complete theory of human rationality could assign truth-conditions to utterances.

(v) So there is no such thing as a semantic theory.

If essentially anything that speakers know about each other and the world can prove relevant to the resolution of context-dependence, then that establishes (iv). As I suggested earlier, however, I do not think the conclusion follows, because I think there is a crucial ambiguity in (iii). I can now explain what that is.

[41]Then again, perhaps not. There is some evidence that even infants are capable of recognizing such intentions (Csibra, 2010).

[42]There is by now excellent evidence that even pre-linguistic children are capable of thinking about objects as such (Carey, 2009).

Expanding on an earlier suggestion due to Burge (1974),[43] Higginbotham (1988; 2002) suggests that the targets of semantic theory should be statements of truth-conditions that are conditioned by claims about the values of contextual parameters. Thus, for example, what a semantic theory should tell us about the sentence "That is a frog" would be something roughly along the following lines:[44]

(1) If u is an utterance of "That is a frog", and if the contained utterance of "that" refers to x, then u is true iff x is a frog.

A semantic theory that issues in claims like (1) is certainly 'about truth-conditions', and the things whose truth-conditions it aims to characterize are utterances. Such a view thus vindicates (i) and is compatible with (ii). Yet it does not commit us to any form of (iii) that might conflict with (iv): The theory does not assign a truth-condition to *any* particular token utterance (even though it does characterize the truth-conditions of utterances), so the fact that it is impossible systematically to assign truth-conditions to individual utterances poses no threat.

It should be clear that this result does not depend upon the details of Higginbotham's proposal (if only because I haven't given any details).[45] The underlying idea is simply that a semantic theory should assign truth-conditions to sentences, but only relative to an assignment of values to contextual parameters. And the crucial point is that semantics need have nothing to say about how contextual parameters get their values: It simply takes for granted that contextual parameters do get values, somehow or other.

So the ambiguity at the heart of the anti-semantic argument can be described as follows. On the one hand, (iii) can be understood as saying that the things whose truth-conditions a semantic theory aims to characterize can only be utterances, never sentences. On this reading, (iii) follows from (i) and (ii) and is true, but it does not conflict with (iv). To get a conflict with (iv), you have to read (iii) as saying that a semantic theory must systematically assign a truth-condition to each *actual* utterance. That is, for each utterance u, it must issue in a categorical statement like:

(2) The utterance u is true iff Kermit is a frog.

So read, however, (iii) does not follow from (i) and (ii), and I see no reason to think it true. As said, on the conception of semantics I would defend, it is of no interest to semantics how the values of contextual parameters are fixed. Semantics only gives us conditionalized T-sentences like (1); semantic theory itself will never[46] put us in a position to discharge the antecedent of the conditional so as to give us something like (2).

[43] Similar suggestions were made independently by Weinstein (1974).

[44] The contrast between "this" and "that" is, of course, not accounted for here, and it is a nice question (on which Brett Sherman has done interesting work) how we might account for it within this sort of approach.

[45] Larson and Segal (1995) develop the proposal in detail.

[46] The only remotely plausible exception would be "I". But Sherman has some nice examples that casts doubt upon whether "I" is an exception.

Still, one might find such a conception unsatisfying. Semantics, one might insist, is born from the thought that utterances carry information (or misinformation) about the world.[47] It is because utterances carry information, and because speakers are pretty good at extracting it, that it is possible for us to communicate about the world *via* language. The goal of semantics is thus to explain how utterances are able to carry information, and the fact that an utterance carries certain information is as objective a fact as you like about it. Any 'semantics' that declines to tell us what information is carried by actual utterances has, therefore, utterly abandoned the very core of what semantics is about.

Well. If semantics is indeed about the information utterances carry, in this sense—if there can be no semantic theory without a theory that tells us what information is carried by which utterances—then, it seems to me, the facts about demonstrative reference that we have reviewed show that semantics is impossible absent a complete theory of human rationality. Or, to put it differently, utterances don't carry information, in the relevant sense, so there is nothing for semantics to be a theory of. But that conception of the goals of semantic theory is optional. A very different conception of semantic theory understands it as responsive to the questions Chomsky has long urged should drive linguistic theory quite generally: What speakers know about their language, how they come to know it, and how that knowledge is put to use (Chomsky, 1965, Ch. 1; 1986, Ch. 1). Semantic theory can be pursued in the same spirit Chomsky would have us pursue syntax or phonology: in an attempt to uncover the cognitive basis of linguistic competence, that is, to make explicit what competent speakers know about their language, on the basis of which they are able to speak and understand it.[48]

This 'cognitive' conception of semantics is clearly compatible with what I am claiming about demonstratives. No doubt it is also consistent with the cognitive conception that, as part of my understanding of "that", I grasp various rules that I use to determine the reference on the basis of contextual cues. But the cognitive conception in no way demands that there should be such rules. If what, *qua* competent speaker, I know about any given utterance of "That is a frog" is given by (1), then what I must do, if I am to assign a truth-condition to a particular utterance of it, is to assign a value to the contained utterance of "that". It is then an empirical question how I do that. If what I do is, as I have argued, simply use whatever I happen to

[47]For example, Soames writes, as if the claim needed no justification: "The central *semantic* fact about language is that it carries information about the world" (Soames, 1988, p. 186, his emphasis). A similar conception is expressed by Salmon (1986). My sense is that it is very widespread, even if it is not always expressed in these terms.

[48]This conception of semantics is of course controversial. The exchange between Higginbotham (1992) and Soames (1992) is a good place to start. The view has roots in the work of Davidson and Dummett, though each of them shies away from a full embrace of it (Heck, 2004). It seems to have been "in the air" in Oxford in the late 1970s. For a sympathetic but ultimately negative assessment, see Martin Davies's paper "Meaning, Structure, and Understanding" (Davies, 1981). Its prominence nowadays is largely due to Higginbotham's influence, both through his own work and through that of his students (of whom I am one). For a detailed elaboration, see the first chapter of *Knowledge of Meaning* (Larson and Segal, 1995).

know to figure out to whom the speaker meant to be referring, then that is of no immediate significance for semantics.[49]

One might worry, however, that this conception of semantics is not available to me. Even conditionalized T-sentences, such as (1), speak of reference and truth. To discharge the antecedent of (1), we would have to establish some such claim as:

(3) The word "that" as it occurs in utterance u refers to Kermit.

thus arriving at (2). But (3) makes a claim about the reference of an uttered demonstrative, and (2) makes a claim about the truth-condition of an utterance. Yet I have denied that uttered demonstratives refer and that utterances have truth-conditions. Are all such claims supposed to be false?

Even if they are all false, I am not sure how bad that is. The view would be that, in using and interpreting language, ordinary speakers ascribe truth-conditions to utterances and reference to uttered demonstratives, which they do not really have. I suppose that would be unfortunate, but, if so, it would be an unfortunate but inevitable fact about the psychology of human speech, comparable to the way some people think about color: Colors do not exist as objective properties of objects, although we do and must perceive them that way, in virtue of how our minds and perceptual systems are designed. Similarly, then: Utterances do not have truth-conditions in their own right, though we do and must hear them that way, due to our how minds and linguistic systems are designed. In short: One could hold a kind of error theory here, and such views are familiar enough.[50]

I confess, however, that I am no fan of error theories, so I would rather avoid any such commitment. The most obvious way to do so is to re-write conditionalized T-sentences like (1) in an explicitly Strawsonian form:

(4) If u is an utterance of "That is a frog", and if the speaker of u uses the contained utterance of "that" to refer to x, then u is true iff x is a frog.

But this would still lead to (2). Perhaps we could try something like:

(5) If u is an utterance of "That is a frog", and if the speaker of u uses the contained utterance of "that" to refer to x, then u is uttered truly iff x is a frog.

But one might independently wonder whether mention of various sorts of linguistic actions should really be made inside semantic theory.[51] We need another approach.

Reference is generally regarded as a binary relation between a token expression and an object. But, while I have denied that words refer, I

[49] I am inclined to go further and suggest that *only* the cognitive conception of semantics can accept both (i) and (ii) without commmiting itself to a form of (iii) that is incompatible with (iv). But I do not really know this to be true. So I shall content myself with the suggestion that the cognitive conception of semantics receives some support from the fact that it does allow us to resolve this antinomy.

[50] Higginbotham (1989) comes close to this view.

[51] For example, if one thinks that language comprehension is, to some extent, modular, then it does not seem plausible that the linguistic module should have access to information about types of linguistic actions.

certainly have not denied that words refer *for speakers* (or, more generally, for conversational participants). So it is open to me to embrace reference as a ternary relation between an expression, an object, and an agent for whom the expression refers to that object: t refers to x for S, or something like that. So, when I deny that uttered demonstratives refer, what I am denying is that reference is a binary relation. The fundamental relation is ternary or, as we might also put it, 'relativized'.

What is it for an expression to refer to an object for a person? I do not know. Indeed, a large part of the point is that, for the purposes of semantic theory, we do not need to know. That is not to say that it is not something we would like to understand, and we shall explore it a bit further shortly. For now, let me say simply that it is supposed to be closely allied to notions like interpretation and understanding. Saying that an expression t refers to an object x for an agent S is thus supposed to be roughly equivalent to saying that S interprets t as referring to x. Or, if one is worried about the use of the verb "to refer": S interprets t as x; or perhaps: S semantically values t as x. The point is that I am using such language to record facts about how S is interpreting a bit of speech, and no allusion to norms about how that speech should be interpreted, or to facts S is trying to get right, is presupposed. We have S, and we have t, and we have how S interprets t, and that is all.

The ternary notion of reference is also of a piece with Strawson's notion of an act of reference, and it allows us, I suggest, better to understand the nature of such acts and the referential intentions with which they are performed. What does it mean to say that S refers to x in uttering a certain demonstrative? It is, I suggest, to say something about how S interprets her own utterance; it is to say that the demonstrative S utters refers to x for her. If so, then we can understand why S can perform such acts at will, since, in planning her utterance, S has already decided how she will interpret it,[52] and, in a case of soliloquy, there will be no more to be said. But in cases that involve communication, the speaker's intention will concern not just how she should interpret her utterance but also how her audience should interpret it. We can understand this intention as concerning what its reference should be for the various members of her audience. We might unify these two ways of regarding the speaker's intention by saying that what she intends is that the demonstrative she is about to utter should refer to x for 'us', where the relevant group includes both the speaker and her intended audience. The speaker's manipulation of the context—pointing or whatever—is then a means she adopts towards this end. Her audience, in turn, must decide how they should interpret the uttered demonstrative, which is simply to say that they must decide what its referent should be for

[52]This formulation is not entirely adequate. Both philosophers and linguists tend to focus on the point of view of the audience, where one is given a sentence and has to interpret it. Very little ever seems to be said about the problem facing the speaker: about how, given that one has decided what to say, one is supposed to find the words to express oneself. I have considered the issue in a few places in my own work (Heck, 2006, 2007), and there are some helpful remarks in a paper by Rumfitt (1995). Maybe the best philosophical discussion I know is in a paper by Pagin (2003). But I do not think we know well enough how to talk about the production of speech, and this lack is surfacing in the text.

them; this decision will be guided by their desire that its referent should be the same for them as it is for the speaker, which we might again express by saying that they need to decide what its referent is for 'us'.

So, as I said, I think the ternary notion of reference is not just available but useful: The phenomena discussed in Section 2 are naturally cast in terms of it.

Back, then, to semantics. I can see no particular reason to suppose that a binary notion of reference must figure in (1) and (2). The phrase "refers" might just mean "refers for me", or "refers for us", or "refers for the speaker", and "true" might be interpreted similarly. Instead of (1) and (2), then, we would have:

(6) The word "that" as it occurs in utterance u refers, for us, to Kermit.

(7) The utterance u is true for us iff Kermit is a frog.

Better still, we might formulate the theory in terms of "refers for x" and "true for x", with the variable filled differently at different times. This would reflect the fact that the theory of truth I tacitly know not only characterizes the truth-conditions that utterances have for me but also the truth-conditions I take those utterances to have for others.[53] Clearly, nothing in the deductive structure of the theory will be affected if we construe it any of these ways. The significance of the theory's output will be affected, to be sure, since the theorems will all involve a relativized notion of truth, but it is not at all clear that this is any cause for concern. Indeed, the discussion in the preceding few paragraphs was meant to suggest that the way I put my semantic knowledge to use in communication precisely involves my figuring out what things mean for me, for us, or for the speaker.

This commitment to a relativized notion *of* truth is not, of course, a commitment to any sort of relativism *about* truth. Classical relativism is a doctrine about the truth of propositions or, in the present framework, interpreted utterances. The relativized truth-predicate in play here is a predicate of *un*interpreted utterances. Saying that a certain utterance is "true for me" just means that it is true given how I interpret it. This would be obvious on a view that took semantics to issue in claims like:

(8) The utterance u means that Kermit is a frog.

The relativized version would then be:

(9) The utterance u means for me that Kermit is a frog.

What is being relativized here is the relation of expression, and the same is true in the case of (6) and (7).

[53] As I have noted elsewhere (Heck, 2006), it is an important but neglected point that how I understand my own words is (excepting the obvious cases) also how I understand those of others.

3.5 Compositionality

A related group of worries focuses on compositionality.

Fodor and Lepore's paper "Out of Context", which has many points of contact with this one, is framed as a discussion of whether context-dependence is *per se* incompatible with compositionality. They write:

> Here's the general idea: by stipulation, a sentence of L is compositional if and only if a (canonical) representation of its linguistic structure encodes all the information that a speaker/hearer of L requires in order to understand it. This means that, if L is compositional, then *having once assigned a linguistic representation to a sentence token, there is no more work for a hearer to do in order to understand it*. ... The upshot is that, if compositionality is assumed, *there is a definite point at which the business of understanding a token of an expression terminates*: it terminates at the assignment of whatever semantic interpretation its linguistic structure determines.
>
> Well, the present objection is that, by these standards, English simply isn't compositional. There would seem to be lots of cases where you need to know more about a sentence token than its linguistic structure in order to interpret it: *you also need to know things about the context of the tokening*. (Fodor and Lepore, 2004, pp. 3–4, their emphasis)

Fodor and Lepore do not mean to endorse this objection.[54] On the contrary, the objection is one they imagine being made to their repeated insistence elsewhere that compositionality is non-negotiable (Fodor and Lepore, 2002). Still, I have quoted them at such length because my own first reaction to this objection is stunned silence. I do not understand why one would suppose that compositionality requires semantic interpretation to be completely determined by linguistic structure. And "by stipulation", nonetheless! Should we really regard compositionality, so understood, as "the key constraint on theories of linguistic content" (Fodor and Lepore, 2004, p. 3)? Is such a strong principle really motivated by the usual sorts of arguments from productivity and systematicity? I think not.[55]

[54]For what it's worth, however, I don't really understand their response to it. One of their central conclusions concerns the priority of thought over language. This suggests that their target might be something like the views discussed in §3.2. But the target they announce is a familiar stalking horse, namely, conceptual pragmatism, and many of their other conclusions are directed at it. What puzzles me, though, is that, so far as I can see, they do not actually defend the strong form of compositionality we are about to discuss. Their own views entail that, in many cases, "you need to know more about a sentence token than its linguistic structure in order to interpret it: you also need to know things about" the speaker's intentions. Why is this supposed to be less of a threat?

[55]There are similar worries that focus on the thesis that semantic interpretation must be modular (Borg, 2004, p. 12), and perhaps this conception of compositionality is motivated by that thought. If so, then the points about to be made apply as well to this view. But the more fundamental point is that, while there may be reason to suppose that *linguistic* interpretation is modular, there is no reason to suppose that *semantic* interpretation must be

The various formal versions of the principle of compositionality are attempts to develop the intuition that the meaning of a complex expression depends upon, and is completely determined by, the meanings of its constituents and they way they are put together. As Szabó (2012) has emphasized—though surely the point goes back to Davidson (1984d, p. 31)—this principle says *nothing* about how the meanings of the constituents are determined. In particular, it leaves entirely open the question whether the meaning of a given constituent might vary from occasion to occasion. Similarly, the principle as stated leaves it open whether the expressions we are discussing are types or tokens. I suppose one might insist that it should be limited to types. But the arguments that motivate compositionality—the arguments from productivity and systematicity—concern the actual use of language, since they concern ordinary speakers' abilities to understand and interpret actually uttered expressions.[56] If so, then it would be natural to suppose that (one acceptable version of) the principle must concern token expressions and so that 'meaning' must *ipso facto* be relativized to occasions of use. What emerges is then the following sort of principle, which is more or less what Szabó (2010) calls 'the compositionality of expression content':

> The content of a complex token expression on the occasion of its utterance is determined by the meanings of its token constituents on the occasion of their utterance and the way the former is constructed from the latter.

Some such principle is very widely endorsed. But the present point is just that, if one is going to suppose that the meaning of a token can depend upon more than just the type it is a token of, then that very supposition suggests a natural version of the principle of compositionality concerning tokens. That makes it hard to take Fodor and Lepore's (interlocutor's) worries seriously.

Still, there is a different worry. In discussing the soldier example, I suggested that the reference one assigns to an uttered demonstrative may depend upon the rest of the sentence in which it occurs: If the uttered sentence is "He must have been up late", the word "he" is most naturally taken to refer to the soldier who fainted; if it is "He's obviously done a lot of

modular, unless one has already decided to limit 'semantics' to context-independent elements of meaning.

[56]Szabó (2010) suggests the contrary: that the argument from productivity, at least, "does not even get off the ground when the relevant notion of meaning is expression content", that is, utterance content. His reason is that knowledge of language, by itself, is insufficient to allow one to understand novel token expressions. But surely that was not the idea. The thought was that one can understand tokens of *some* types one has never encountered before, whereas it is *not* true that one can understand tokens of just *any* type one has never encountered before. Why only some types, then? Because a significant, though incomplete, contribution is made by something one knows about the sentence uttered.

Granted, this does not establish the compositionality of expression content. It is possible that the meaning one assigns the expression should be derived from the meaning one assigns to the uttered sentence via a pragmatic process similar to implicature, so that the meaning assigned to the expression might have little to do with that assigned to the sentence. But that just means that we need additional arguments, perhaps of the sort Szabó gives, perhaps of the sort King and Stanley (2004, p. 128) give. Arguments to be given below, about why we need a notion of 'what is said', are also relevant here.

practicing", then to the soldier performing the drill; if "He's obviously done a lot of practicing to faint on cue", then to the one who fainted; and so on and so forth. Isn't this a straightforward violation of compositionality? It seems that the reference of the demonstrative depends upon the meanings of other parts of the sentence to which it is not, in any plausible sense, structurally related.

This objection really is guilty of conflating the metaphysical and epistemic senses of "determine". To see this, note that essentially the same points could be made about cases of lexical ambiguity. If Bob says, "Joe jumped off the bank and swam across the river", then one may not realize that it is the word that means *river bank* that has been uttered until one hears the rest of the sentence.[57] This certainly does not show that which word 'bank' has been uttered depends upon the meanings of other parts of the sentence. What is true is that *one determines* which word has been uttered using what one knows about other parts of the uttered sentence; this has no tendency to show that the meaning of this particular token of the word 'bank' *is determined* by anything about other parts of the uttered sentence.[58]

The same goes for demonstratives. What is true is that, in many cases, *one determines* what value to assign an uttered demonstrative on the basis of one's understanding of the rest of sentence in which it occurs. This is quite compatible with its being the case that the meaning of the uttered sentence *is determined* by the meanings of its parts and the way they are combined.

None of this, I should emphasize, is intended as an argument for compositionality. I am only arguing that reasonable forms of the principle of compositionality are no less plausible if the conception of context-dependence I have been defending is accepted.

3.6 Semantics and Pragmatics

Another batch of objections concerns what the view I am defending implies about the distinction between semantics and pragmatics.

A first objection is that any view that gives central place to the speaker's intentions, as mine to some extent does, will contravene the distinction between speaker's reference and semantic reference. Recall the following sort of example, due to Kripke (1977, pp. 263–4). Out across the way we see a man working in his yard, and one of us asks, "What is Jones doing?" Answer: "Raking the leaves". But it isn't Jones; it's Smith. Then in some sense we have referred to Smith; but the name "Jones" never refers to Smith but only to Jones. So the *semantic* referent of the expression "Jones" was Jones, but the *speaker's* referent was Smith.

[57] Nowadays, it is commonly held that there are really two words 'bank' that are pronounced alike. But the remarks in the text do not depend upon this claim and can easily be modified if it is not accepted.

[58] Much the same can be said about structural ambiguity and even about merely sonic similarities. Suppose one cannot quite tell whether it was a 'g' or a 'c' that filled the blank when Martha said, "I had my -oat cleaned, and they shrunk it". Then one's understanding of the second conjunct and general knowledge of the world will help one decide what word Martha uttered and so what she said.

The other day, a friend of mine was over, and I saw one of our cats lying on the sofa. I said to him, "That cat loves to have his belly rubbed". I said this because I thought it was my cat Snarfy, who really does love to have his belly rubbed. In fact, however, it was his sister Lily, who enjoys a belly rub from time to time but isn't nearly as enthusiastic about them as Snarfy is. So I intended to refer to Snarfy, and yet it seems clear that I did not. Shouldn't we say, then, that the *semantic* referent was Lily, but the *speaker's* referent was Snarfy? If so, isn't the semantic referent determined, not by the speaker's intentions, but rather by something like contextual cues?

No. To see this, we need to be more specific about the intentions with which I spoke in the example just given. The first point to note is that I clearly *did* intend to refer to the cat on the sofa. Contrary to what I believed, that cat was Lily, but it does not follow that I did not intend to refer to her: Intention is intensional. Nor does the fact that I intended to refer to the cat on the sofa imply that I did not also intend to refer to Snarfy. I did. If I did not think that the cat on the sofa was Snarfy, then I would not have said what I did. So I had two intentions, to refer to the cat on the sofa and to refer to Snarfy, and these intentions were linked: I intended to refer to Snarfy *by* referring to the cat on the sofa; I thought I could do this because I thought that the cat on the sofa *was* Snarfy. This fact about the structure of my intentions surfaced in my expectations about how I would be interpreted. My friend did not yet know Snarfy's name, so I did not expect him to interpret my utterance of "that cat" as referring to Snarfy but rather to the cat on the sofa. (Interpretation is intensional, too.) Or, better: I expected him to interpret my utterance as referring to *that* cat, that is, to interpret the utterance in terms of a demonstrative thought. So what I expected him to conclude, if he took me at my word, was that *that* cat likes to have his belly rubbed.[59]

We should distinguish, then, between what we might call *proximal* and *distal* referential intentions.[60] Exactly how this distinction should be drawn is not a question we can discuss in detail here, so I shall have to leave it at this intuitive level.[61] The point is that, once we make such a distinction, we shall want to say that what my utterance of a demonstrative refers to, even for me, is determined by my *proximal* intentions. So, in this case, even for me, the uttered demonstrative referred to Lily, not to Snarfy, because my proximal intention was to refer to *that* cat, and that cat is Lily. (Reference is *ex*tensional.)[62]

[59]As Kaplan notes in "Afterthoughts" (Kaplan, 1989a, §II), essentially the same point applies to his Carnap–Agnew case: Kaplan does not just intend to refer to the picture of Carnap but also to the picture behind him, and he intends to refer to the former *by* referring to the latter.

[60]This distinction is similar in spirit to one drawn by Bach (1992).

[61]I have discussed related matters elsewhere (Heck, 2006, pp. 82ff), as has Rumfitt (1995, pp. 850ff).

[62]With what right do I so casually speak of our thoughts about "*that* cat"? How do we know what "*that* cat" refers to? Isn't our topic demonstratives? Yes, but our topic here is *language*, that is, the use of demonstrative *expressions*; our topic is not demonstrative *thought*. The two are undoubtedly related. In many uses of demonstratives, the speaker will think of the object to which she intends to refer demonstratively: That is, she will intend to refer to *that*

Still, one might think I face an even worse problem. The worry is that, once speaker's intentions are granted central importance, they cannot but assume overriding importance. Fodor and Lepore, whose views are in many ways similar to mine, deny that "... the right interpretation of an utterance depends on the linguistic conventions that determine what the speaker 'strictly and literally says'" (Fodor and Lepore, 2004, p. 12).[63] What they mean is apparent from a footnote.

> Suppose someone says "It's raining here" believing, wrongly, that he is in Pittsburgh. Then the right interpretation for the hearer to impose on the utterance is *it's raining in Pittsburgh*, not *it's raining here*. To insist on the 'literal' meaning would be pointless and pedantic, not to say uncharitable. (Fodor and Lepore, 2004, fn. 25)

The thought is obviously that the speaker intends to say that it is raining in Pittsburgh, and Fodor and Lepore want to insist that, since the speaker's intention dominates, then that *is* what he said.[64] But that just seems wrong.

And it is wrong. Suppose this person has just been asked if he'd like to go for a picnic. Then whether he looked out the window or checked the weather online, the fact that he believes he is in Pittsburgh is of no significance. What he has said is that it is raining where he is; that is the only thing relevant to the question at issue. But perhaps Fodor and Lepore have a different sort of case in mind. Suppose that what has been asked is whether the Pirates game will be cancelled.[65] Our confused speaker looks out the window and says, "It's raining here". Then there is a sense, no doubt, that what the speaker *means* is that it is raining in Pittsburgh. But we can explain that intuition in familiar terms: The speaker's remark would be irrelevant unless he thought that he was in Pittsburgh, in which case the fact that it was raining where he was would imply that it was raining in Pittsburgh. But nonetheless, what he has *said* is that it is raining where he is.

What allows us to distinguish what the speaker said from what the speaker meant is the fact that the speaker used the word "here", which, as a matter of its meaning—of what it means in his language, of what it means to him—cannot be used to refer to Pittsburgh simply because that is where

---

thing, and she will expect her audience to interpret her by forming demonstrative thoughts of their own. But I see no reason to suppose that non-anaphoric uses of demonstratives must always involve demonstrative thoughts in this way, as Evans (1982, §9.1) sometimes seems to suggest. Cases of 'deferred ostension' are clear counter-examples unless they can be set aside. But it is hard to see why one would want to set them aside unless one was antecedently committed to the view that demonstrative reference requires demonstrative thought.

[63]What Fodor and Lepore strictly and literally say is that certain considerations due to Evans do not establish this claim. But they give every impression that they want to deny it.

[64]One might suggest that Fodor and Lepore simply mean to deny that how one ought to interpret the speaker is determined by what he strictly and literally said. But that has been obvious since Grice, if not before. That is: If they are not just pointing out that what the speaker means may differ from what the speaker said, then the scare quotes around "strictly and literally says" must be there because they are skeptical about the very notion. As, in some ways, am I, though there are different ways to be skeptical about it.

[65]The Pirates are Pittsburgh's professional baseball team. Baseball games can be postponed due to rain.

one thinks one is. It would be too strong to say that it must be used to refer to the place where one is. One can say, pointing at a location on a map, "Here is where the treasure is". Precisely what restrictions the use of "here" does impose is therefore unclear. But the essential point is that the use of "here" *must* impose *some* restrictions on what one can say by uttering it. Its imposing those restrictions is what makes it a word with meaning rather than so much nonsense.

As obvious as this point is, it often seems to be overlooked. Example after example has been given over the last several years to show that one can say many different things by uttering even very simple sentences, such as "That ink is blue":[66] One might mean that the ink appears blue in the bottle, or that it will appear blue upon drying, or yet other things. But little notice ever seems to be taken of how few things might be said by uttering this sentence, that is, of how many things there are that *cannot* be said by uttering it. For example, under no circumstances can one say, by uttering the sentence "That ink is blue", that the existence of a measurable cardinal implies that there are non-constructible sets. And that is a very good thing. If one could say absolutely anything by uttering absolutely any sentence, then the job of interpretation would be infinitely harder than it is.[67] The point is just a variant of one Chomsky made famous: If too many possible languages were an option for the child, then it would be all but impossible for children to acquire languages under the conditions they actually do. Similarly: If there were too many possible interpretations of an uttered sentence, then it would be all but impossible for us to interpret it under the actual conditions we do.

So far as I can see, therefore, absolutely everyone has to acknowledge that the words one uses, even if they do not determine what one says in uttering them, nonetheless sharply *constrain* what one can say by uttering them. That, as I said, is what makes them words rather than the verbal equivalent of a gesture. The only question worth discussing is therefore *how* the words one uses constrain what can be said by uttering them.

Note carefully how that was phrased: The words one uses restrict what one can *say* by uttering them. It should be clear that only "say" is appropriate here. The words one uses do not constrain what one can, in Grice's sense, *mean* by uttering them, since what one means includes what one implicates, and it is easy to see that, given any proposition and any sentence, there are circumstances in which one can implicate that proposition by uttering that sentence.[68] Of course, the words one uses will make it easier or harder to mean this or that, but the notion of constraint to which I'm appealing is different: It is the one at work when one says, for example, that the sentence "He saw John in the mirror" does not have as a reading that John saw himself

[66]This particular example is due to Travis (2000). Similar examples have been given by Sperber and Wilson (1995), Recanati (2004), and many others.

[67]Maitra (2007, §5) makes some similar remarks, though her topic is the problem of shared content, to be discussed in Section 4.

[68]A boring way to get this result is to assume that some conditional $p \to q$ is sufficiently salient to the conversational participants, where q is what one wants to implicate and p is what the utterance of the sentence will semantically express. But it's not hard to come up with more interesting examples.

in the mirror. We have to understand this as: The sentence cannot be used to *say* that John saw himself in the mirror. We have to understand it that way because the sentence can, of course, be used to *mean* anything you like, given appropriate circumstances.

This observation allows me to answer a question some readers will have been wanting to ask for some time, namely, why we need a notion of 'what is said' here at all or, better, why we need a 'fully propositional' notion of what is said. The alternative would be to try to make do with (i) a notion of what is said that is essentially context-independent and is characterized in terms of the meaning of a sentence and (ii) a context-bound notion of what a speaker means. The former notion is not 'fully propositional', because sentences as such do not express propositions but only what are sometimes called 'gappy' propositions or 'proposition radicals'. For example, the sentence "That is a frog" would express a gappy proposition we might indicate as: x is a frog. Communication would then be understood as consisting, in the general case, in the expression of gappy propositions, through which process complete propositions are meant.[69]

One objection to this sort of view is that it seems to conflate very different sorts of pragmatic processes. It just feels wrong to lump the determination of the reference of a demonstrative in with the calculation of implicatures. And, clearly, the calculation of implicatures will in many cases depend upon the resolution of context-dependence: What is implicated will depend, e.g., upon which thing has been said to be a frog. But the issue is complicated, since there are also cases in which one can implicate without expressing a complete proposition at all. Indeed, it is easy to construct cases—one can start with Strawon's famous example, "That's a fine red one"—in which the implicature turns precisely on the fact that no proposition is expressed. So it is not true that implicature always requires the resolution of context-dependence and so must turn on what is (propositionally) said.

But there is a better answer, namely, that the most natural explication of the notion of what a sentence means is in terms of how it constrains the linguistic acts that can be performed by uttering it, and that these acts must be characterized not in terms of what is meant but in terms of what is *said*. As I understand the sentence "That is a frog", for example, its meaning makes a certain range of linguistic acts available to me: By uttering this sentence, I can *say* of any object I choose that it is a frog; I cannot *say* anything else. But, as noted, it is only what I can *say* that is so constrained, so the notion of what is said underlies the notion of sentence-meaning.[70]

It is worth illustrating these themes with a more complex sort of example. Consider the much discussed phenomenon of quantifier domain restriction. As mentioned earlier, a sentence like "Everyone is on the bus" can apparently be used to say different things on different occasions. There are several accounts of why. One such account supposes that relevant material is elided; another supposes that the speaker can freely 'enrich' the proposition ex-

[69]Bach (2004b) is the best-known proponent of this sort of view.

[70]The point can be put more colorfully this way: If one wants to speak of the meaning of a sentence as a 'gappy' proposition, then that will make good sense only if some special role is assigned to the result of filling the gaps.

pressed by the words themselves; yet a third supposes that there is no context-dependence here at all and that the phenomenon is pragmatic. But there has been little effort by the proponents of these accounts to address the critical question why sentences containing quantifiers have the very restricted range of meanings they do. Consider, for example, these three sentences:

(10) Everyone who read a book skipped some pages.

(11) Everyone who read skipped some pages.

(12) Everyone skipped some pages.

It seems to me that (10) has the reading: Everyone who read a book skipped some pages *in each of the books she read*. Neither (11) nor (12) can have that reading, however, even if it is clear, in the case of (11), that only books were being read and, in the case of (12), that the only people under discussion are those who read a book.

How should this be explained? It is not easy to see what any of the accounts mentioned so far might have to offer. If the restriction in place in (12) is to be explained by elision, enrichment, or pragmatics, then it is far from clear why (12) could not be an elided form of (10), why (12) could not be enriched to (10), or why the pragmatic processes could not extract (10) from (12). I do not say that no such explanation could be given. I simply point out that no such explanation has been given or even, so far as I am aware, seriously attempted.[71] It is a significant advantage of the view defended by Stanley and Szabó (2007),[72] according to which quantifiers are associated with covert arguments, that it offers, and indeed is to some extent motivated by, precisely such examples.[73]

But whatever the right explanation, every reasonable view must significantly constrain what one can say by uttering "Everyone is on the bus". One way or another, such a view must ensure that the speaker's freedom is limited, so far as "everyone" is concerned, to fixing the domain over which that quantifier will range. (There are similar things to be said about "the bus".) No matter how fervent, the speaker's intentions cannot make "everyone" mean anything other than *everyone in group G*, but can only affect which group *G* is at issue. Maybe that isn't quite right. Maybe the speaker's freedom is more or less extensive than I have just suggested. But it doesn't matter. The point is that there simply have to be some such constraints on the speaker's freedom, whether they lie in the syntax, in the lexicon, in the

[71] Hall (2008, §4) does discuss this sort of issue, in an attempt to defend free enrichment. By her own account, however, her discussion is just a beginning, and the locality constraint she mentions does not seem to address the example just presented.

[72] Fodor and Lepore (2004, p. 5) discuss this sort of view briefly, saying only that they "don't know of anything principled" to be said in favor of it. Perhaps what I am saying here counts.

[73] Stanley and Szabó do not discuss this sort of example, but a suitable story falls directly out of their account. Very roughly, when (10) has the reading we are discussing, it has roughly the structure of "Everyone who read a book skipped some pages *in it*". Thus, (10) has this reading because one of its structures is that of a donkey sentence, and donkey sentences typically have universal readings. But (11) cannot have this reading for much the same reason that "Everyone who read skipped some pages in it" is ungrammatical.

rules governing enrichment, or in some as yet unimagined place. We would not be talking about language otherwise.

Giving the speaker's intentions a central role does not, therefore, imply giving the speaker's intentions the sort of over-riding importance that would obliterate the distinction between semantics and pragmatics, between what is said and what is meant.[74] As I have emphasized, everyone must agree that the meaning of the word "everyone" constrains what can be said by uttering it. As before, it is what one can *say* that is so constrained, not what one can, in Grice's sense, *mean*; so that distinction survives intact.[75] But if we draw the distinction between what is said and what is meant in these terms, then we need not try to connect it to what is more or less objective, nor to what is more or less rule-governed. In particular, we do *not* need to suppose that context itself somehow determines what is said in such cases. If, for example, the right thing to say about "everyone" is that it must always mean: everyone in group G, so that the locus of context-dependence is the domain for the quantifier, then what I have been arguing is that the domain is not (and does not need to be) fixed 'by context', in the sense that objective facts about the circumstances in which an utterance is made will fix this parameter.[76] Rather, in uttering the sentence "Everyone is on the bus", the speaker will (proximally) intend to speak of the members of some group G, and her utterance will, for her, express the proposition that everyone in that group is on the bus. Her audience will then have the task of determining which group this is, and successful communication will depend upon their performing this task correctly. There is no need for an 'objective standard' at which we are all aiming—though we might now see the desire for one as an understandable over-reaction to the need for language to constrain what can be said.

Other cases can be handled similarly. Consider gradable adjectives, like "tall", whose contribution to the significance of a given utterance seems to vary with the occasion. One popular story, developed in work by Kennedy (1999) and others, has it (again very roughly) that such adjectives are associated with an appropriate scale—a scale of height, in the case of "tall"—and that such an adjective always means something like: is higher (or lower) than some cut-off point δ. Again, the details of this story may be wrong, but the spirit of the story must be right: The speaker's freedom must be limited and, if we side with Kennedy for the moment, then it is limited to selecting the cut-off point. No matter how fervently one might intend otherwise, then, one cannot utter "John is tall" and thereby say that John is a fish (though, of course, one could *mean* that John is a fish). What I have been arguing then concerns how δ is fixed. I claim that it is not fixed 'by context' but by the speaker and her audience, each for herself, and that communication depends

[74]It may be that part of the reason it can seem as if I would be committed to such a conclusion is the fact that the linguistic meanings of demonstratives place so few constraints on what they can be used to say.

[75]And so, in that sense, Grice is absolutely right that "what someone has said [is] closely related to the conventional meaning of the words... he has uttered" (Grice, 1989, p. 25).

[76]Independent arguments could of course be given concerning this case. Gauker (1997) discusses some relevant examples. See note 31 for discussion of his treatment of them.

upon convergence.

The case of color is much harder, but we don't understand color, anyway. Still, some significant work has been done on the semantics of color words (Szabó, 2001; Reimer, 2002; Rothschild and Segal, 2009; Kennedy and McNally, 2010), and I see no reason to despair. As I have said, there have to be constraints on the range of things one can mean by "blue". Whatever those constraints are, the view view for which I've been arguing concerns how the meaning the word has on a particular occasion is, and is not, determined.

So the worry was that the emphasis on speakers' intentions—my 'admission' that pragmatic processes affect what is said—somehow undermines the very notion of 'what is said' and leads us down the road to relevance theory, truth-conditional pragmatics, or what have you. But it should be clear by now, I hope, that there is nothing in the position for which I have argued that should incline one toward such views. Quite the contrary, in fact. Such views, it seems to me, can only seem less attractive once the 'semantic' perspective to which they are opposed has been freed from any association with the view that 'what is said' must somehow be determined by rule.

4 Closing

Whether or not uttered demonstratives 'objectively' refer, speakers who utter demonstratives will intend to speak about particular objects, and their audiences will interpret their utterances as being about various objects, with successful communication requiring (at least) agreement. The question I have been asking is: Once this broadly psychological story has been told, what role is left for the 'objective' referent of the uttered demonstrative? I see none. We do not need it to explain successful communication (§3.1); we do not need it to make sense of natural language semantics generally (§3.4) or, in particular, of compositionality (§3.5); and we do not need it to respect the distinction between semantics and pragmatics (§3.6).

Now, as I mentioned at the outset, and as will have become clear in the last section, the discussion of demonstratives has been intended as a window onto context-dependence quite generally. The overall picture is that, where *E* is an expression whose meaning may vary from occasion to occasion, then *E*'s context-independent meaning will fix a range of meanings it can have on a given occasion of utterance, and the selection of a particular meaning is accomplished not 'by context' but, rather, by the speaker and her audience, each for herself, with successful communcation requiring convergence.

This picture is almost present in a paper by Glanzberg. He argues that the contextual parameters associated with gradable adjectives, in particular, require what he calls an "indirect metasemantics", which he describes in these terms:

> A contextual parameter with an indirect metasemantics must be set by the various pieces of information context provides, but context does not simply hand us a value for such a parameter, nor does it hand us a uniform rule for computing the value from a specific piece of contextual information. Rather, a range of

> contextual information and computational rules must be taken into account and weighed in working out the value from context. (Glanzberg, 2007, p. 19)
>
> What fixes [such] values will be complicated combinations of such factors as what is salient in the environment, speakers' intentions, hearers' intentions, coordinating intentions, linguistic meaning, general principles governing context, discourse structure, etc. From these resources, values will have to be computed. . . . [S]uch computations will often allow for accommodation and negotiation. (Glanzberg, 2007, p. 25)

I could not agree more—except that Glanzberg never quite recognizes the full implications of this way of thinking.[77] Who is performing these computations? Well, who is accommodating and negotiating? The speaker and her audience, obviously. So these computations are not ones the theorist performs to determine 'what is said', but ones the speaker and her audience perform in an effort to understand one another. The better way to put Glanzberg's point would therefore be: Such contextual parameters *do not have* a 'metasemantics' in the usual sense; *there is no story to be told* about how context fixes their values, because it doesn't. Rather, values are assigned to the contextual parameters by the speaker and her audience, and the process through which they determine what values to assign is not governed by "a uniform rule" but is shaped by a common goal: successful communication, which requires them to agree. It is unsurprising that accomodation and negotation should be central to this enterprise, and it is equally unsurprising that so many "factors" should be involved. As we have seen, there is, in fact, no limit to the information on which the speaker and her audience may draw as they attempt to converge.

But now one might wonder just how successful this attempt typically is. Consider again the case of "tall". How likely is it that, in ordinary sorts of cases, such as when Peggy says to Sue, "Bill is tall", the speaker and her audience really *do* manage to converge on a value for the parameter, on a cut-off point? It does not seem particularly likely. If not, then won't we have to say that, in a wide range of quite ordinary cases, in which we might have thought we were communicating successfully, we actually aren't?

Before I address that question, let me raise another. Forget about communication. Is it really plausible that we always have a precise cut-off point in mind when we utter the word "tall"? Mabye not. If so, however, then, according to me, the speaker in such cases does not express any specific proposition, since, according to me, the only thing there is to determine that proposition is the speaker's intention,[78] and I have just conceded that the

[77] Glanzberg (2007, p. 25) also contrasts this sort of case with the case of demonstratives, about which he remarks that "[t]here are roughly two options". Obviously, I disagree, but his case against relativism is only strengthened if it turns out that even demonstratives do not fit the model of context-dependence that the relativist assumes. I am happy to acknowledge, too, that the process of fixing (or, perhaps better, agreeing on) a value for a parameter might be very different depending upon the parameter.

[78] Careful! I am not claiming here that the speaker's intention determines what proposition is expressed, in some objective sense, but only what proposition the *speaker* expresses.

speaker's intention is insufficiently precise to determine such a proposition. And the same presumably goes for the speaker's audience. Since they probably do not have any specific cut-off in mind, they do not understand the speaker to be expressing any specific proposition.

These issues are undoubtedly worth extended discussion, and they are not going to get it here. But let me mention four points.

First, we can see in such considerations another reason one might want to appeal to a mechanism that fixes values for contextual parameters on the basis of objective features of context: By fixing a cut-off point for us, perhaps context can determine both what proposition the speaker expresses and what proposition she is understood by the audience to have expressed, and do so in such a way as nearly to guarantee agreement.[79]

Second, it has not been widely appreciated[80] that this problem, which we may call the 'problem of shared content', threatens every view except the one just mentioned: that context objectively resolves context-dependence. In particular, so-called 'minimalist' views (Borg, 2004; Cappelen and Lepore, 2005), according to which a sentence like "The ink is blue" actually does express some fixed proposition (*modulo* tense), are just as threatened by it as my view is. This is because, on the minimalist view, this fixed proposition, although 'semantically expressed' by the speaker's words, is almost never what she communicates (and has, indeed, no very clear role to play, though that is a different issue). The fact that the speaker's words express some fixed proposition thus does not help us understand how successful communication occurs; pragmatic processes play at least as big a role for the minimalist as they do for me, and as they do for people who like gappy propositions. So, while the problem is real, it is solvable in principle, since (or so I have argued here) the view that context objectively resolves context-dependence is untenable.

Third, unless you are an epistemicist about vagueness (Williamson, 1994), it should on reflection not seem very plausible that context *does* fix a definite cut-off point for a given utterance of "tall". Indeed, cases like "tall" seem to me to pose a serious problem for epistemicism, for just this reason. But, however that may be, it does not seem unreasonable to suppose that the propositions speakers ordinarily express using "tall" are, in some sense we would like better to understand, imprecise. And while it would be a mistake to confuse this sort of imprecision with vagueness, my own suspicion is that the two are not unrelated (Heck, 2003).[81]

Finally, it is generally assumed that successful communication requires,

[79]The view to which such considerations lead is thus very much Gauker's: His point is largely that the right conception here is a sort of externalism according to which it is not just the contents of our utterances but also the contents of our *thoughts* that are fixed by contextual factors (Gauker, 2003). For what it's worth, my own view is that such a conception of thought, which is explored in some detail by Gross (2005), is incoherent.

[80]A clear exception is Maitra (2007).

[81]Nor am I alone. The sort of 'dynamic' picture of vagueness one finds in Tappenden (1993; 1995), according to which 'precisification' is a very real part of our use of vague terms, is very much of a piece with the view I have just expressed. Some of the ideas expressed by Shapiro (2006) are in a similar vein, too, despite the fact that his view has 'contextualist' aspects that I would reject.

in the relevant sorts of cases, that the audience should come to entertain the very proposition that the speaker has expressed. In fact, however, this is not at all clear. I have argued elsewhere that the case of demonstratives is itself a counterexample (Heck, 2002). In a way, that case is very different, because my view is not that communication might succeed despite disagreement about reference. Rather, I regard the 'propositions' speakers believe and express as fine-grained, individuated not just by reference but also by (something like) sense, and I think speakers and their audiences do not have to agree about *sense* to communicate successfully. Still, the case is suggestive. Perhaps successful communication does not require the speaker and her audience to interpret her utterance as expressing the very same proposition, but only as expressing 'similar' propositions. The problem, of course, is to say what "similar" means. But it has long been noted in the literature on vagueness that lack of a sharp cut-off will have no practical effect so long as its precise location does not affect the categorization of the objects presently under discussion.[82]

I suspect that this idea can be developed in such a way as to allow us to articulate an appropriate notion of similarity. But that is different from saying I have so developed it, so further discussion of this issue will have to be deferred to another occasion.[83]

References

Almog, J., Perry, J., and Wettstein, H., eds. (1989). *Themes from Kaplan*. New York, Oxford University Press.

Bach, K. (1992). 'Intentions and demonstrations', *Analysis* 52: 140–6.

—— (2004a). 'Context *ex Machina*', in Szabó 2004, 15–44.

—— (2004b). 'Minding the gap', in C. Bianchi (ed.), *The Semantics/Pragmatics Distinction*. Chicago, University of Chicago Press, 27–43.

[82]That is: The extension of "tall", restricted to the objects presently under discussion, is not affected by the precise location of the cut-off. This idea in some sense goes back to Fine's (1975) work on supervaluationism, but it has been emphasized more recently by Fara (2000) and Kennedy (2007).

[83]Thanks to Alexi Burgess and Brett Sherman for helpful comments that improved the paper. Thanks also to Tony Corsentino and to Steven Gross (as well as to Brett) for lengthy discussions of these issues that significantly shaped my thinking about them, and to Jason Stanley, with whom I've been discussing this topic for a long time. The position developed here is meant to be broadly compatible with the work he has done on context-dependence (Stanley, 2007).

Versions of this material were presented at the University of Connecticut, in April 2008; at the Arché Center for Logic, Language, Metaphysics, and Epistemology, at the University of St Andrews, in June 2008; to a discussion group at Brown University, in March 2010; at the University of California at Berkeley, in April 2010; at the Northern Institute of Philosophy, at the University of Aberdeen, in June 2010; and at the University of Massachusetts, in March 2011. Thanks to everyone who attended these lectures for their reactions, but especially to Kent Bach, Derek Bowman, Herman Cappelen, Hannah Ginsborg, Ishani Maitra, Robert May, John McFarlane, Randall Rose, Brian Weatherson, Sam Wheeler, Crispin Wright, Seth Yalcin, and Michael Young for comments that did much to improve the paper. Thanks to Crispin for welcome encouragement, too.

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

Burge, T. (1974). 'Demonstrative constructions, reference, and truth', *Journal of Philosophy* 71: 205–23.

Cappelen, H. and Lepore, E. (1997). 'On an alleged connection between indirect speech and the theory of meaning', *Mind and Language* 12: 278–96.

—— (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford, Blackwell.

Carey, S. (2009). *The Origin of Concepts*. New York, Oxford University Press.

Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge MA, MIT Press.

—— (1986). *Knowledge of Language: Its Origin, Nature, and Use*. Westport CT, Praeger.

Csibra, G. (2010). 'Recognizing communicative intentions in infancy', *Mind and Language* 25: 141–68.

Davidson, D. (1984a). *Inquiries Into Truth and Interpretation*. Oxford, Clarendon Press.

—— (1984b). 'Radical interpretation', in Davidson 1984a, 125–139.

—— (1984c). 'Thought and talk', in Davidson 1984a, 155–70.

—— (1984d). 'Truth and meaning', in Davidson 1984a, 17–36.

Davies, M. (1981). 'Meaning, structure and understanding', *Synthese* 48: 135–61.

Evans, G. (1982). *The Varieties of Reference*, McDowell, J., ed. Oxford, Clarendon Press.

Fara, D. G. (2000). 'Shifting sands: An interest-relative theory of vagueness', *Philosophical Topics* 28: 45–81. Published under the name 'Delia Graff'.

Fine, K. (1975). 'Vagueness, truth, and logic', *Synthese* 30: 265–300.

Fodor, J. and Lepore, E. (2004). 'Out of context', *Proceedings and Addresses of the American Philosophical Association* 78: 77–94.

—— (2002). *The Compositionality Papers*. Oxford, Clarendon Press.

Gauker, C. (1997). 'Domain of discourse', *Mind* 106: 1–32.

—— (2003). *Words Without Meaning*. Cambridge MA, MIT Press.

Glanzberg, M. (2007). 'Context, content, and relativism', *Philosophical Studies* 136: 1–29.

Grice, H. P. (1989). 'Logic and conversation', in *Studies in the Ways of Words*. Cambridge MA, Harvard University Press, 22–40.

Gross, S. (2005). 'Context-sensitive truth-theoretic accounts of semantic competence', *Mind and Language* 20: 68–102.

Hall, A. (2008). 'Free enrichment or hidden indexicals?', *Mind and Language* 23: 426–56.

Heck, R. G. (2002). 'Do demonstratives have senses?', *Philosophers' Imprint*, 2. `http://www.philosophersimprint.org/002002/`.

—— (2003). 'Semantic accounts of vagueness', in J. C. Beall (ed.), *Liars and Heaps: New Essays on Paradox*. Oxford, Oxford University Press, 106–27.

—— (2004). 'What is compositionality?', *Manuscript*.

—— (2005). 'Reason and language', in C. MacDonald and G. MacDonald (eds.), *McDowell and His Critics*. Oxford, Blackwells, 22–45.

—— (2006). 'Idiolects', in A. Byrne and J. J. Thomson (eds.), *Content and Modality*. Oxford, Oxford University Press, 61–92.

—— (2007). 'Meaning and truth-conditions', in D. Greimann and G. Siegwart (eds.), *Truth and Speech Acts: Studies in the Philosophy of Language*. New York, Routledge, 349–76.

Higginbotham, J. (1988). 'Contexts, models, and meanings: A note on the data of semantics', in R. Kempson (ed.), *Mental Representations: The Interface Between Language and Reality*. Cambridge, Cambridge University Press, 29–48.

—— (1989). 'Knowledge of reference', in A. George (ed.), *Reflections on Chomsky*. Oxford, Basil Blackwell, 153–74.

—— (1992). 'Truth and understanding', *Philosophical Studies* 65: 3–16.

—— (2002). 'Competence with demonstratives', *Philosophical Perspectives* 16: 1–16.

Kaplan, D. (1978). 'Dthat', in P. Cole (ed.), *Pragmatics*. New York, Academic Publishers, 221–43.

—— (1989a). 'Afterthoughts', in Almog et al. 1989, 565–614.

—— (1989b). 'Demonstratives', in Almog et al. 1989, 481–563.

Katz, J. and Fodor, J. (1963). 'The structure of a semantic theory', *Language* 39: 170–210.

Kennedy, C. (1999). 'Gradable adjectives denote measure functions, not partial functions', *Studies in the Linguistic Sciences* 29: 65–80.

—— (2007). 'Vagueness and grammar: The semantics of relative and absolute gradable adjectives', *Linguistics and Philosophy* 30: 1–45.

Kennedy, C. and McNally, L. (2010). 'Color, context, and compositionality', *Synthese* 174: 79–98.

King, J. C. (2001). *Complex Demonstratives: A Quantificational Account*. Cambridge MA, MIT Press.

King, J. C. and Stanley, J. (2004). 'Semantics, pragmatics, and the role of semantic content', in Szabó 2004, 111–64.

Kripke, S. (1977). 'Speaker's reference and semantic reference', *Midwest Studies in Philosophy* 2: 255–76.

Larson, R. and Segal, G. (1995). *Knowledge of Meaning*. Cambridge MA, MIT Press.

Lewis, D. (1998). 'Index, context, and content', in *Papers in Philosophical Logic*. Cambridge, Cambridge University Press, 21–45.

Maitra, I. (2007). 'How and why to be a moderate contextualist', in Preyer and Peter 2007, 112–32.

Montague, R. (1974). 'Pragmatics', in R. Thomason (ed.), *Formal Philosophy: Selected Papers of Richard Montague*. New Haven, Yale University Press, 95–118.

Mount, A. (2008). 'Intentions, gestures, and salience in ordinary and deferred demonstrative reference', *Mind and Language* 23: 145–64.

Pagin, P. (2003). 'Communication and strong compositionality', *Journal of Philosophical Logic* 32: 287–322.

Pagin, P. and Pelletier, F. J. (2007). 'Content, context, and composition', in Preyer and Peter 2007, 25–62.

Preyer, G. and Peter, G., eds. (2007). *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*. Oxford, Oxford University Press.

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

Reimer, M. (1991). 'Demonstratives, demonstrations, and demonstrata', *Philosophical Studies* 63: 187–202.

—— (1992). 'Three views of demonstrative reference', *Synthese* 93: 373–402.

—— (2002). 'Do adjectives conform to compositionality?', *Philosophical Perspectives* 16: 183–98.

Rescorla, M. (2006). 'Review of Christopher Gauker, *Words without Meaning*', *Philosophical Review* 155: 121–4.

Rothschild, D. and Segal, G. (2009). 'Indexical predicates', *Mind and Language* 24: 467–93.

Rumfitt, I. (1995). 'Truth-conditions and communication', *Mind* 104: 827–62.

Salmon, N. (1986). *Frege's Puzzle*. Cambridge MA, MIT Press.

Shapiro, S. (2006). *Vagueness in Context*. Oxford, Clarendon Press.

Soames, S. (1988). 'Semantics and semantic competence', in S. Schiffer and S. Stelle (eds.), *Cognition and Representation*. Boulder CO, Westview Press, 185–207.

—— (1992). 'Truth, meaning, and understanding', *Philosophical Studies* 65: 17–35.

Sperber, D. and Wilson, D. (1995). *Relevance: Communication and Cognition*, 2d edition. Oxford, Blackwell.

Stalnaker, R. (1999a). 'Assertion', in Stalnaker 1999b, 78–95.

—— (1999b). *Context and Content: Essays on Intentionality in Speech and Thought*. New York, Oxford University Press.

—— (1999c). 'On the representation of context', in Stalnaker 1999b, 96–113.

Stanley, J. (2007). *Language in Context*. Oxford, Oxford University Press.

Stanley, J. and Szabó, Z. G. (2007). 'On quantifier domain restriction', in Stanley 2007, 69–110.

Strawson, P. F. (1950). 'On referring', *Mind* 59: 320–44.

—— (1952). *Introduction to Logical Theory*. London, Methuen.

—— (1971). 'Meaning and truth', in *Logico-Linguistic Papers*. London, Methuen, 170–89.

Szabó, Z. G. (2001). 'Adjectives in context', in R. M. Harnish and I. Kenesei (eds.), *Perspectives on Semantics, Pragmatics, and Discourse*. Amsterdam, John Benjamins, 119–46.

—— (2010). 'The determination of content', *Philosophical Studies* 148: 253–72.

—— (2012). 'The case for compositionality', in E. M. W. Hinzen and M. Werning (eds.), *The Oxford Handbook of Compositionality*. Oxford, Oxford University Press. Forthcoming.

Szabó, Z. G., ed. (2004). *Semantics versus Pragmatics*. Oxford, Oxford University Press.

Tappenden, J. (1993). 'The liar and sorities paradoxes: Toward a unified treatment', *Journal of Philosophy* 90: 551–77.

—— (1995). 'Remarks on vagueness and a dynamic conception of language', *Southern Journal of Philosophy: Spindel Conference Supplement* 33: 193–201.

Travis, C. (2000). *Unshadowed Thought*. Cambridge MA, Harvard University Press.

Weinstein, S. (1974). 'Truth and demonstratives', *Noûs* 8: 179–84.

Wiggins, D. (1997). 'Languages as social objects', *Philosophy* 72: 499–524.

Williamson, T. (1994). *Vagueness*. London, Routledge.

Wittgenstein, L. (1968). *Philosophical Investigations*, 3d edition, tr. by G. E. M. Anscombe. Oxford, Basil Blackwell.