

Nonindexical Contextualism*

John MacFarlane (UC Berkeley)

December 16, 2007

Abstract

Philosophers on all sides of the contextualism debates have had an overly narrow conception of what semantic context sensitivity could be. They have conflated context sensitivity (dependence of *truth* or *extension* on features of context) with indexicality (dependence of *content* on features of context). As a result of this conflation, proponents of contextualism have taken arguments that establish only context sensitivity to establish indexicality, while opponents of contextualism have taken arguments against indexicality to be arguments against context sensitivity. Once these concepts are carefully pulled apart, it becomes clear that there is conceptual space in semantic theory for *nonindexical* forms of contextualism that have many advantages over the usual indexical forms.

1 Introduction

It is common for philosophers to call an expression “context-sensitive” just in case its content varies with the context in which it is used:

To say that e is context sensitive is to say that its contribution to the propositions expressed by utterances of sentences containing e varies from context to context. (Cappelen and Lepore, 2005, 146)

*This paper was presented at the Rutgers Semantics Workshop in September, 2005 and the UCLA Philosophy Colloquium in October, 2005. For helpful feedback I am indebted to audiences at both talks, my commentators at Rutgers (Delia Graff Fara and Zoltan Szabó-Gendler), the Bay Area Philosophy of Language Discussion Group (especially Kent Bach, Chris Barker, Michael Glanzberg, Mike Martin, Line Mikkelsen, and Michael Nelson), my Spring 2005 graduate seminar on context sensitivity, Anonymous, Stewart Cohen, Stephen Neale, and Jason Stanley.

A sentence is context-sensitive if and only if it expresses different propositions relative to different contexts of use. (Stanley, 2005, 16)

To say that vague predicates are context sensitive is to say that they are indexical. While the semantic content of an indexical varies from one context of utterance to another, its meaning does not. (Soames, 2002, 245)

In what follows, I will argue that we should understand context sensitivity more broadly, as dependence of *extension* on context. It is possible for an expression to be semantically context-sensitive, in this sense, even if it has the same content at every context of use.¹

The point is not merely that we are associating the *word* “context-sensitive” with the wrong concept—after all, we are free to use technical terms in whatever way is most useful. The problem is that philosophers are using only one word where there are two distinct concepts. Because these concepts are not distinguished terminologically, they are often conflated. As a result of this conflation, arguments on both sides of the contextualist debates are guilty of equivocation, and a distinctive kind of contextualist view does not even appear as a conceptual possibility.

As a first step towards conceptual clarity, I suggest that we use the word “indexicality” for the dependence of content on context, reserving “context sensitivity” for the dependence of extension on context. Thus,

- (1) An expression is *indexical* iff its *content* at a context depends on features of the context.²
- (2) An expression is *context-sensitive* iff its *extension* at a context depends on features of the context.

The sense of “indexical” defined by (1) is quite broad. It does not distinguish between different ways in which the content of an expression might depend on the context. Sometimes “indexicality” is used in a narrower sense to cover just some of these ways (Stanley, 2000, 411). If you like, call the sense defined by (1) “broad indexicality.”

¹Henceforth I will omit “semantically” before “context-sensitive.” There are many kinds of non-semantic context sensitivity, but they are not at issue here. See Stainton, forthcoming.

²I’ll assume here that the content of a sentence at a context is a proposition, and that the content of a subsentential expression is the contribution it makes to the content of sentences containing it.

These generic notions of indexicality and context sensitivity are too indiscriminating for many purposes. Typically we are interested in whether an expression is semantically sensitive to *certain aspects* of the context of use: for example, the time of the context, or the epistemic standards that are in play at that context. So it will be useful to work with parameterized versions of the notions defined in (1) and (2). Where P is a feature of contexts,

- (3) An expression is *P-indexical* iff its content at a context depends on the P of that context.
- (4) An expression is *P-context-sensitive* iff its extension at a context depends on the P of that context.

Thus, someone who thinks that

- (5) It is night in New York City

expresses different propositions depending on the time of use holds that it is time-indexical. Someone who thinks that

- (6) At the moment when he finished his most famous article, Moore knew that he had hands

has different truth values depending on the epistemic standards in play at the context of use holds that it is epistemic-standard-context-sensitive. And so on.

Note that “depends” in (4) has causal/explanatory force. To show that the truth value of S depends on feature P , it is not enough just to find two contexts that differ with respect to P and relative to which S has different truth values. For the difference in truth values may be due to other differences between these contexts. It is generally not possible to find pairs of contexts that differ in respect P without differing in many other ways as well.

It is easy to see that indexicality does not entail context sensitivity, and P -indexicality does not entail P -context sensitivity. The sentence “Tomorrow comes after today,” for example, is time-indexical—its content at a context depends on the time of the context—but not time-context-sensitive, since it has the same extension (True) at every context of use. In what follows, I will argue that the converse entailments also fail: an expression can be context-sensitive but not indexical, and for many interesting choices of P , an expression can be P -context-sensitive without being P -indexical. So P -context-sensitivity and P -indexicality are independent. As we will see, however, some central arguments both for and against contextualism conflate

P -indexicality with P -context-sensitivity, and are thus invalid due to equivocation. Distinguishing the two notions reveals attractive forms of contextualism that are otherwise invisible. As an alternative to holding (with standard contextualism) that an expression is P -indexical, one can take the view that it is P -context-sensitive but not P -indexical. Nonindexical contextualism is immune to many of the objections that have been leveled against standard (indexical) forms of contextualism. It faces some unique objections of its own, but these are not as forceful as they might first appear.

2 A motivating example

There's an old debate about whether the truth of propositions can vary with time. *Temporalists* hold that it can. According to temporalism, the sentence "Socrates is sitting" expresses the same proposition at 2 PM as it does at 3 PM: the proposition that Socrates is sitting. This proposition has truth values only relative to times: it may be true relative to 2 PM, false relative to 3 PM. The opposite viewpoint is taken by *eternalists*, who hold that propositions have their truth values eternally. According to eternalism, the sentence "Socrates is sitting" expresses different propositions at different times. At 1400 UTC on January 1, 412 BC, it expresses the proposition that Socrates is sitting at 1400 UTC on January 1, 412 BC. At 1500 UTC, it expresses a different proposition, the proposition that Socrates is sitting at 1500 UTC on January 1, 412 BC.

Temporalists and eternalists do not disagree about the truth values of occurrences of tensed sentences. They agree, for example, that an occurrence of "Socrates is sitting" at time t is true just in case Socrates is sitting at t . They just disagree about the mechanism by which the sentence acquires these truth conditions. On the eternalist's view, the sentence varies in truth value across times because it expresses different propositions at different times, and these propositions have different (eternal) truth values. On the temporalist's view, the sentence varies in truth value across times because the (single) proposition the sentence expresses at all these times has different truth values relative to different times of evaluation.

To get clearer about how temporalists and eternalists can agree about the truth of sentences while disagreeing about the truth of propositions, it may help to look at what David Kaplan (a temporalist) says about the relation between sentence and proposition truth:

If c is a context, then an occurrence of ϕ in c is true iff the content expressed by ϕ in this context is true when evaluated

with respect to the circumstance of the context. (Kaplan, 1989, 522)

Notice that the context c plays two distinct roles in determining sentence truth. First, it helps determine which proposition is expressed by the sentence. I'll call this the *content-determinative* role of context. Second, it tells us at which circumstance of evaluation we should evaluate this proposition to get a truth value for the sentence in context. Since circumstances for Kaplan are world/time pairs, it tells us *which* world and time to look at: the world and time of the context of utterance. I'll call this the *circumstance-determinative* role of context.

We can now describe the difference between the temporalist and the eternalist as follows. Both agree that the truth value of a tensed sentence depends on the time of the context of use. That is, they agree that such sentences are time-context-sensitive. But where the eternalist takes the time of the context to play a content-determinative role, the temporalist takes it to play a circumstance-determinative role. Thus, the eternalist, but not the temporalist, takes tensed sentences to be time-indexical. The temporalist's position, according to which tensed sentences are time-context-sensitive but not time-indexical, is an example of *nonindexical contextualism*.

If you found yourself objecting to my definition of "context sensitivity" as dependence of extension on context, ask yourself this question: Do you really think that temporalists should deny that tensed sentences are context-sensitive? Shouldn't they rather say that tensed sentences are context-sensitive, even though they are not indexical in the broad sense? (You may well reject temporalism,³ but nothing in my argument hangs on accepting it. What matters is that it's an intelligible semantic theory, one that can find a place in orthodox semantic models like the one found in Kaplan (1989). It needn't be the *best* theory.)

3 Contingency is a kind of context sensitivity

Temporalism provides a clear example of context sensitivity that is not attributable to indexicality, but we need not go so far afield. Take any contingent, eternalized sentence without indexicals—say,

(7) On May 18, 1993, Bill Clinton got his hair cut.

³For useful discussions of the issues, see King (2003) and Richard (2003).

This sentence expresses the same proposition at every context—even contexts located at different possible worlds. Nonetheless, its truth value varies with the context. Uttered in the actual world, it expresses a truth; uttered in countless other possible worlds, it would express a falsehood. In this respect it is context-sensitive.⁴

David Lewis expresses this point well:

When truth-in-English depends on matters of fact, that is called *contingency*. When it depends on features of context, that is called *indexicality*. But need we distinguish? ... It is a feature of any context, actual or otherwise, that its world is one where matters of contingent fact are a certain way. Just as truth-in-English may depend on the time of the context, or the speaker, or the standards of precision, or the salience relations, so likewise may it depend on the world of the context. Contingency is a kind of indexicality. (Lewis 1980 in Lewis 1998, 25)

When he says that contingency is “a kind of indexicality,” Lewis is using “indexicality” the way I am recommending we use “context sensitivity,” to mean dependence of sentence truth on features of context.⁵ He is certainly not using it to mean contextual variation of *content*. Indeed, a main point of his paper is that semantics need not concern itself with propositional contents at all. Nor is he using it to mean dependence of truth value on the *index*. When he talks of the dependence of “truth-in-English” on “the world of the context,” he is talking not about the three-place relation “*S* is true at context *c* and index *i*,” but about the two-place relation “*S* is true at context *c*,” defined as follows:

Let us say that sentence *s* is true at context *c* iff *s* is true at *c* and the index of the context *c*. (Lewis 1980, 31; cf. Kaplan (1989, 522), quoted above)

Just as, for Kaplan, the time of the context affects the truth of sentences by helping determine which circumstance of evaluation is “the circumstance of the context,” so, for Lewis, the world of the context affects the truth of sentences by helping determine which index is “the index of the context.”

⁴Appreciating this point does not require being a modal realist. It just requires a willingness to consider merely counterfactual uses of sentences.

⁵Lewis’s use of “contingency” is also a bit nonstandard. We typically now use “contingent” as a predicate of *propositions*, not sentences. As Kaplan (1989) made vivid, a sentence that expresses a contingent proposition can nonetheless be true relative to every context of use: witness “I am here now” (cf. Lewis (1998, 29)).

So, in contrast to the contemporary trend of operating only with the *narrower* notion of context sensitivity—what I called “indexicality” above—Lewis uses only the broader notion. It is not surprising, then, that contemporary semanticists have found this passage confusing.⁶ Its basic insight, however, is worth preserving: sentences that express the same proposition at every context may yet be context sensitive, in the sense that their truth depends on features of the context.

You might say: “Surely a notion of context sensitivity on which (7) counts as context-sensitive is too all-encompassing to be theoretically useful!” But the generic notion of context sensitivity is a limit case. In practice, the parameterized notions, like time-context-sensitivity and epistemic-standard-context-sensitivity, are of more interest. Let us see, then, what can be gained by distinguishing epistemic-standard-indexicality from epistemic-standard-context-sensitivity. I focus on contextual sensitivity to epistemic standards primarily because the debates about it are so familiar, but the lessons learned should be applicable in many other domains as well.

4 Nonindexical epistemic contextualism

Standard contextualist views take “know” to be indexical, in my broad sense. On these views, knowledge-attributing sentences express different propositions at different contexts of use.⁷ Is there room for a *nonindexical* contex-

⁶Stanley (2005, 132 n. 1) argues that Lewis’s claim is really a point about restricted quantifier domains: “For Lewis, if quantifiers are read unrestrictedly, sentences have their truth-values necessarily. So, according to modal realism, contingency does result from a kind of context-sensitivity—namely, quantifier domain restriction. Non-modal realists cannot appeal to Lewis’s claim as a defense of an implausibly broad usage of ‘context-sensitivity’.” But only some contingent sentences involve quantifiers: (7), for example, does not, and its contingency has nothing to do with quantifier domains. Nor is there any mention of contextual determination of quantifier domains in this passage—or anywhere else in Lewis’s article.

⁷DeRose (1996, 194), Schiffer (1996, 318), Stanley (2004, 119), and Schaffer (2004, 73) all characterize contextualism as a view about the dependence of *content* on context. DeRose (1992, 914), DeRose (2002, 167), and Feldman (2001, 62) talk, less committally, of the dependence of *truth conditions* on context. There are at least six different things that might be meant by “truth condition” here: (i) function (in the mathematician’s extensional sense) from contexts to truth values, (ii) rule for determining truth values based on features of context, (iii) function from circumstances of evaluation to truth values, (iv) rule for determining truth values based on features of circumstances of evaluation, (v) function from possible worlds (and perhaps times) to truth values, (vi) rule for determining truth values based on features of possible worlds (and perhaps times). All of these notions can come apart. Because it is rarely clear which of them is meant by “truth condition,” that

tualism about “know”—a view that would stand to standard contextualism as temporalism stands to eternalism?

Using temporalism as a model, we can see what such a view would look like. Temporalists can have time-context-sensitivity without time-indexicality because they take proposition truth to be relative to a world/time pair. So let’s take proposition truth to be relative to a world/epistemic standard pair. We’ll call such a pair a “circumstance of evaluation.”

We want this to be a *nonindexical* contextualism, so let’s say that the word “know” expresses the same relation at every context of use: the relation a subject S stands in to a proposition p and a time t just in case S knows that p at t . This relation, like others, has an *intension*: a function from circumstances of evaluation (world/epistemic standard pairs) to extensions. We can specify this intension roughly as follows: a subject S stands in the knowing relation to p and t at (w, e) iff p is true at (w, e) and S is in a strong enough epistemic position at w and t with respect to p to satisfy the standard e .

Since “know” always expresses the same relation, on this view, it is not indexical. To see how it can be *context-sensitive*, we have to consider the relation between sentence truth (relative to a context of use) and proposition truth (relative to a circumstance of evaluation). Here we can just follow Kaplan:

If c is a context, then an occurrence of ϕ in c is true iff the content expressed by ϕ in this context is true when evaluated with respect to the circumstance of the context. (Kaplan, 1989, 522)

We just need to specify what counts as “the circumstance of the context” when circumstances are world/epistemic standard pairs:

(w, e) is the circumstance of the context c iff w is the world of c and e is the epistemic standard in play at c .

On this view, the truth values of sentences containing “know” depend on the epistemic standard in play at the context of use, not because this standard affects which proposition is expressed, but because it helps determine which circumstance of evaluation to look at in deciding whether these sentences are true or false at the context. The semantics is structurally just like the temporalist’s semantics discussed above.

term is best avoided altogether.

It is worth making explicit how this nonindexical form of contextualism differs from other, more familiar views about the semantics of “know.” It differs from strict invariantism in taking the truth of knowledge-attributing sentences to depend on contextually variable epistemic standards. It differs from standard indexical forms of contextualism in taking “know” to express the same relation at every context. And it differs from subject-sensitive invariantism (Hawthorne, 2004; Stanley, 2005) in taking the epistemic standard relevant to the truth of knowledge-attributing sentences to be fixed directly by the context of use, not by the subject’s situation at the time and world of evaluation. Thus the nonindexical contextualist can agree with strict and subject-sensitive invariantists that “know” expresses the same relation at every context of use, with contextualists and subject-sensitive invariantists that “know” is sensitive to varying epistemic standards, and with strict invariantists and contextualists that the standards for “knowing” do not vary with shifts in the time and world of evaluation.⁸

Finally, nonindexical contextualism differs from relativism (Richard, 2004; MacFarlane, 2005a) in taking the epistemic standard parameter to be initialized by the context of use, rather than the context of assessment. For this reason, it may be attractive to those who find the notion of assessment-relative truth too much to swallow. It has many of the same advantages as relativism over standard contextualist and invariantist views, but requires less of a departure from standard semantic frameworks. (A potential disadvantage will be discussed later.)

Kompa’s contextualism

Has such a view about knowledge attributions ever been proposed and defended? Not with full attention to the semantic details. But this view would make good sense of the (otherwise puzzling) position taken by Nikola Kompa in her article “The Context Sensitivity of Knowledge Ascriptions” (Kompa,

⁸Subject-sensitive invariantism predicts that sentences like

- (8) Jim knows that he will be in London next year, but he didn’t know this when he was considering whether to buy insurance (even though he believed it and had exactly the same grounds he has now), and
- (9) Sarah doesn’t know that she will get a job when she graduates, but she would have known this had she not been responsible for her mother’s medical bills (in which case less would have been at stake).

should be true. This is generally considered to be a bad result, though there is disagreement about how bad it is. See DeRose (2000), Hawthorne (2004, 166-8), MacFarlane (2005a, 202), Stanley (2005, 106-14).

2002). Unlike most other writers, who describe contextualism as the view that the *contents* or *truth conditions* of knowledge-attributing sentences vary with context (see note 7), Kompa describes contextualism as the view that “the *truth value* of a given knowledge ascription depends in a certain way on the context of utterance of that very ascription” (80, emphasis added). That is, she defines contextualism as the view that “know” is context-sensitive, *not* as the view that it is indexical. She then notes that “[i]t is quite common among contextualists to account for the context dependence of knowledge ascriptions by pointing to the alleged indexicality of the word ‘know’” (85). After dismissing such views on the basis of the arguments from speaker’s knowledge in Schiffer (1996), she proposes that we understand the context sensitivity of “know” (as well as “flat”, “good”, “tall”, and many other words) not as indexicality but as “unspecificity.” An unspecific predicate expresses the same property at every context of use, but “what counts as having this property might vary from context to context” (88). As a result, different utterances expressing the same proposition may have different truth values at different contexts:

...an unspecific utterance is true or false, as the case may be, only relative to the imposed standard. The standard in turn is determined by contextual features like the speaker’s and hearer’s presuppositions, interests, intentions, their conversational goals &c. So a truth condition of an unspecific utterance could be roughly stated as follows, where ‘... is F’ be an unspecific predicate:

An utterance of “X is F” [sic] is true iff X meets the contextually relevant standard for F-ness (87-8)

Kompa concludes: “All of this points to the fact that an expression may well be context sensitive without being an indexical, ambiguous, elliptical or vague expression” (92).

Everything Kompa says fits the model of nonindexical contextualism about “know” I described in the last section perfectly. However, she isn’t explicit enough about the semantic details for us to be sure that this is what she intended.⁹

⁹Perhaps Peter Ludlow has in mind something like nonindexical contextualism when he raises the possibility that “‘knows’ is a context-sensitive predicate without an L-marked position for standard or degree of knowledge and that knowledge reports have no operators representing standards of knowledge.” He notes: “It could still be the case that ‘knows’ is a context-sensitive predicate. Here I am thinking that ‘knows’ could work like tense mor-

5 Conflating context sensitivity with indexicality

I must confess that when I first read Kompa's article, I found it unintelligible. How can she say that the truth values of utterances of knowledge-attributing sentences depend on contextually supplied epistemic standards, I wondered, if these standards play no role in determining the propositional contents of these utterances? I thought she was just confused. But it was I who was blind. Nonindexical contextualism was invisible to me. There was simply no room in my conceptual space for such a position, because I was conflating context sensitivity with indexicality.

I don't think I was alone in my blindness: the same blind spot is evident throughout the literature on epistemic contextualism. Once we have distinguished context sensitivity from indexicality, we can find them being conflated in arguments on both sides of the debates. Proponents of contextualism take arguments for semantic context sensitivity to be arguments for indexicality, and opponents take arguments against indexicality to be arguments against context sensitivity. It's worth looking at a couple of examples.

Indirect reports

Opponents of contextualism (including Hawthorne (2004) and Cappelen and Lepore (2005)) make much of the fact that we tend to report attributions of knowledge homophonically, even when they were made in contexts where different epistemic standards were relevant. Even if such reports may sometimes be a bit misleading, we don't hear them as literally *false*. Witness the oddity of the following dialogue:

Sam: I know that my car is in the driveway.

(A bit later, in a context with different epistemic standards in play...)

Barry: Sam said/asserted/believed that he knew that his car was in the driveway.

Janet: No, he didn't!

The most straightforward explanation of why Janet's reply seems so wrong is that it is false: Barry's report of Sam's claim (or of the belief it expressed) is correct. But the contextualist cannot concede this. Barry's report represents Sam as having said/asserted/believed the proposition Barry expresses

phemes for A-theorists—they are context sensitive but there is no explicit argument place for times" (Ludlow, 2005, 27).

with the words “he knew that his car was in the driveway.” But on the contextualist account, this proposition is not the same as the one Sam expressed with the words “I know that my car is in the driveway.” The word “know” expresses a different relation in Barry’s mouth than it did in Sam’s mouth. So Barry’s report must misrepresent the content of Sam’s claim—which is exactly what Janet is saying.

Of course, the contextualist may be able to give a less straightforward explanation of why Janet’s reply seems wrong. Perhaps she is implicating a falsehood by asserting a truth, or perhaps we (and other ordinary speakers) suffer from a kind of semantic blindness that prevents us from seeing (without philosophical assistance) that “know” is indexical. I think that the prospects of giving a compelling explanation along these lines are dim, but I won’t argue that here.¹⁰

An alternative response is available to contextualists who take “know” to be “incomplete” in some way, rather than indexical in the narrow sense. These contextualists hold that the propositions expressed by sentences containing “know” vary with context, not because the semantic value of “know” itself is fixed by the epistemic standards in play at the context, but because “know” is completed differently at different contexts (for example, with different contrast classes, on Schaffer’s approach).¹¹ Their model for the context sensitivity of “know” is not “I”, but “local.” Thus they might say that, just as Barry in Berkeley can report Sam’s utterance (in Santa Fe) of

(10) I’m going to a local bar

by saying

(11) Sam said he was going to a local bar,

meaning *local-to-him*, so Barry can report Sam’s utterance of

(12) I know that my car is in the driveway

by saying

(13) Sam said that he knew his car was in the driveway,

meaning *knew-by-the-standards-operative-at-his-context*. Perhaps Barry has discretion to complete “know” the same way Sam did (even though his epistemic context is very different), or perhaps he can make the completion anaphoric on “Sam”.

¹⁰See MacFarlane (2005a).

¹¹Note that these views still count as “indexical” contextualisms in my broad sense.

The problem with this approach (as Stanley (2004, 138) points out) is that we don't seem to be able to complete "know" differently at different places in the same discourse (especially sentence), the way we can with "local." If we could, then we would be able to hear an utterance of

- (14) Sam knows that his car is in his driveway, but he doesn't know that it hasn't been stolen.

as saying something true (with the two occurrences of "know" completed differently). But it's doubtful that we can: such sentences seem almost contradictory. As Stanley points out, contextualists have typically not wanted to allow true readings of (14); DeRose, in particular, attaches great importance to the fact that his contextualism renders such "Abominable Conjunctions" false (DeRose, 1995, 27-9). This gives contextualists a reason to take "know" to be an "automatic" indexical, on the model of "I". But then they cannot handle the indirect report data, or so it is alleged.

Note, however, that the argument from indirect reports cuts only against *indexical* forms of contextualism. *Nonindexical* contextualists hold that the proposition Sam expresses by saying "I know that my car is in the driveway" is the same as the proposition Barry later expresses, in a different epistemic setting, by saying "Sam knew that his car was in the driveway," so that the oddity of Janet's response can be explained in the obvious way: it is just false. Stanley's observation about (14) is explained as well; the context of use determines a single epistemic standard that is relevant to the evaluation of both occurrences of "know". Nonetheless, those who use the argument from indirect reports typically present it as a general argument against the view that "know" is semantically context sensitive. Indexicality is conflated with context sensitivity, and nonindexical contextualism becomes invisible.¹²

Context-shifting arguments

I have suggested that by embracing nonindexical contextualism, contextualists might insulate themselves against what is commonly taken to be a telling argument against contextualism. But contextualists tend to conflate indexicality with context sensitivity just as much as their critics do. Indeed, the main arguments they use to establish their position—so-called "context-shifting arguments"¹³—*depend* on this equivocation.

¹²The criticism of contextualist solutions to the sorites paradox in Stanley (2003) has similar limitations, though in fairness Stanley's main target (Scott Soames) explicitly advocates an indexical form of contextualism.

¹³The terminology comes from Cappelen and Lepore (2005).

A context-shifting argument (CSA) exhibits a single sentence (S) and asks us to consider whether S would be truly uttered in two different contexts of use, C_1 and C_2 . We are supposed to have the intuition that S would be truly uttered in C_1 but not in C_2 , and this is supposed to show that S has a different content in C_1 than in C_2 . Here's a nice example of this kind of reasoning:

Since in [context] O I'm ascribing knowledge to the same Henry of whom I deny knowledge in [context] S , to hold that both of my utterances are true *requires* the 'contextualist' idea that there's a change in the content of 'knows' between these cases. (DeRose 1996, 194, emphasis added)

Arguments of this form are absolutely central to the contextualism debates. Interestingly, hardly anyone on either side of these debates questions the *validity* of CSAs.¹⁴ Opponents of contextualism typically focus on the *premises* of these arguments, denying that the intuitions about truth and falsity elicited by the cases are accurate guides to the literal truth and falsity of the sentences in context. But once nonindexical contextualism is in view as a live option, it should be clear that CSAs are not even valid. They can show, at best, that the target sentence is P -context-sensitive—that its truth value relative to a context of use depends on some feature P of the context (say, the epistemic standards in play).¹⁵ But this is not enough to show that the sentence is P -indexical—that the proposition it expresses depends on feature P of the context of use. As we have seen, a sentence can be P -context-sensitive without being P -indexical.

It's worth diagnosing the flaw in CSAs in a little more detail. We describe two utterances of the same sentence, S , one in context C_1 , the other in context C_2 . We then observe that intuitively one utterance is true, while the other is not. What this means is that these utterances express propositions that are true and untrue, respectively, at the circumstances of their respective contexts. Assuming our intuitions about the truth of these utterances are accurate, we can conclude that

¹⁴Cappelen and Lepore (2005) are an admirable exception. Although I think their criticism of CSAs from the standpoint of "speech act pluralism" is illuminating, I think it is striking that they, too, fail to see the more fundamental flaw with CSAs that I will describe below.

¹⁵The reason for the qualification "at best" is that some CSAs do not distinguish between features of the subject's situation and features of the context of use, e.g. because the sentences they use are first-person present-tense. The truth intuitions in such cases can be explained just as well by a subject-sensitive invariantism. A successful CSA must ensure that the *only* difference between the cases are differences in the relevant features of the context of use.

- (15) at C_1 , S expresses a proposition that is true at the circumstance of C_1 ,
and
- (16) at C_2 , S expresses a proposition that is not true at the circumstance
of C_2 .

We *cannot* conclude, however, that the proposition S expresses at C_1 is different from the proposition S expresses at C_2 . For if the circumstance of C_1 is different from the circumstance of C_2 , our two utterances of S might diverge in truth value even while expressing the same proposition.

For example, suppose S is the sentence “Bush won the US election in 2004,” and suppose that the world of C_1 is distinct from the world of C_2 . Then an utterance of S in C_1 could diverge in truth value from an utterance of S in C_2 , not because different propositions are expressed, but simply because the circumstances of the two contexts are different. (Say, Bush won in the world of C_1 , but lost in the world of C_2 .)

In order to establish that different propositions are expressed in the two contexts, a CSA would need an additional premise:

- (17) The circumstance of C_1 is the same as the circumstance of C_2 .

I have yet to see this premise made explicit in a CSA. Granted, it is fairly easy to secure in orthodox frameworks, where a circumstance of evaluation is just a possible world. In such frameworks, (17) is true whenever the two contexts are situated at the same world. So the user of a CSA has only to describe contexts that take place at the same world and differ only in other ways, and the CSA *will* establish that S expresses different propositions at C_1 and C_2 . But if circumstances of evaluation are, say, world/epistemic-standard pairs, as we contemplated earlier, premise (17) will be very difficult to secure. It is no longer sufficient to ensure that C_1 and C_2 are situated at the same world; we must also make sure that the same epistemic standard is in play in these contexts. But if we do that, we no longer get premises (15) and (16) for our target sentences. Thus, unless they are accompanied by some principled argument against including an epistemic standards parameter in circumstances of evaluation (on which more later), CSAs are powerless to establish the indexicality of “know”.

To summarize: CSAs *do* show that “know” is context-sensitive (provided the premises can be defended), but not (without further premises) that “know” is indexical, as most contextualists have held. On the other hand, indirect discourse arguments tell against the claim that “know” is indexical, but don’t touch the more general claim that “know” is semantically context-sensitive.

Thus nonindexical contextualism looks like a good option for those who are impressed by the usual arguments against contextualism but still feel the force of CSAs.

6 Nonstandard parameters

One way to block interesting forms of nonindexical contextualism would be to give some principled argument for limiting circumstances of evaluation to “standard” parameters: worlds and perhaps times.

Incompleteness

Sometimes the following consideration is invoked. Propositions are supposed to be the contents of propositional attitudes. But if we specify the content of someone’s attitude in a way that leaves its accuracy undetermined, we have not given its complete content. Thus, for example, if we don’t know whether the accuracy of Sam’s belief that it is 0° C depends on the temperature in London on Tuesday or the temperature in Paris on Wednesday, we don’t yet have the full story about what it is that Sam believes.

This is a familiar enough line of thought,¹⁶ but it proves too much. For surely the accuracy of *any* contingent belief depends on features of the world in which the believer is situated (the world of the context of use). Even if we say that Sam believes that it is 0° C at the base of the Eiffel Tower at noon local time on February 22, 2005—making the time and place fully explicit—we have left it undetermined whether Sam’s belief is accurate. That depends on whether he is in a world in which Paris (or its counterpart) is having a particularly mild February or one in which Paris (or its counterpart) is having a cold February. I have used modal realist language here, but nothing depends on that. All that is necessary is a willingness to consider whether Sam’s belief would have been accurate had the temperature been different.

One might respond to these considerations by bringing the world of the context of use into the *content* of Sam’s thought. Intuitively, though, Sam could have had a thought with the very same content even if the world had been very different. Our ordinary ways of individuating thought contents do not support making the world of the context of use part of the content, except in exceptional circumstances. Moreover, bringing the world of the

¹⁶See e.g. Frege (1979, 134–5).

context into the content of Sam's thought would make this content a necessary truth about this possible world, rather than a contingent truth about the weather in Paris. We should not say, then, that Sam's thought is *about* the world of use. It is not *about* any particular world. Acknowledging the fact that it depends for its truth on the world of use, we may adopt John Perry's terminology and say that it *concerns* the world of use (Perry, 1986).

The objection from "incompleteness" may be motivated, in part, by an appreciation of the fact that the truth predicate we use in ordinary speech is monadic. We don't characterize claims as "true-in-*w*," or as "true-in-*w*-at-*t*-on-*s*," but as "true" (simpliciter). But this no more shows that propositional truth is not relative to parameters than the fact that we normally say it's "3 PM," and not "3 PM Pacific Daylight Time," shows that the time of day is not relative to a time zone. The monadic predicate "true" is just another piece of vocabulary whose intension we can characterize using the relation of truth at a circumstance of evaluation: its extension at a circumstance of evaluation *e* is just the set of propositions that are true-at-*e*.

Once we accept the relativity of propositional truth to worlds, we have accepted a kind of "incompleteness." We have accepted the idea that both the content of an assertion or belief and its context must be taken into account in assessing it for accuracy. The question is just *which* features of context must be taken into account, once the content has been specified. Just about everyone will say the world of the context; temporalists will add the time of the context; and an epistemic nonindexical contextualist will add the contextually relevant epistemic standards.

It is worth noting that the considerations that philosophers have brought to bear against temporalism—largely considerations about our practices in reporting attitudes (Richard, 1980, 2003)—don't have much force against nonindexical epistemic contextualism. For, as we have seen, our practices in reporting knowledge claims do not support the idea that epistemic standards are part of the content. If anything, then, epistemic standards are *more* plausible than times as parameters of propositional truth.

Conservatism

One might fall back on the objection that allowing "nonstandard" parameters would constitute a radical break from hallowed tradition. But in this case there *isn't* a hallowed tradition. Granted, it has become something of an orthodoxy in the last decade or two to take circumstances of evaluations to be just worlds. But let's not forget that Kaplan used a time parameter and left his options open for other kinds of parameters:

A circumstance will usually include a possible state or history of the world, a time, and perhaps other features as well. The amount of information we require from a circumstance is linked to the degree of specificity of contents, and thus to the kinds of operators in the language. (Kaplan, 1989, 502)

If we wish to isolate location and regard it as a feature of possible circumstances we can introduce locational operators: ‘Two miles north it is the case that’, etc. Such operators can be iterated and can be mixed with modal and temporal operators. However, to make such operators interesting we must have contents which are locationally neutral. (504)

These words come from one of the founding documents of the tradition in formal semantics now regarded as orthodox.

Operators

A more principled argument for rejecting an epistemic standards parameter can be found in Stanley (2005, 147-52):

- (a) We should only countenance a parameter of circumstances if there is an operator that shifts it.
- (b) There is no operator that shifts epistemic standards.
- (c) Therefore, we should not countenance an epistemic standards parameter.

I won’t question premise (b) here. Stanley (2005) points out that “strictly speaking,” “by high standards,” and “by the standards of chemistry” are not plausibly understood as epistemic-standards-shifting operators, since they seem to have nonredundant effects on sentences that don’t contain epistemic vocabulary (151-2). Whether there are better candidates—for example, “on any reasonable standard for knowledge” or “by the epistemic standards appropriate to the law courts”—is, I think, an open question. So I will leave this (broadly empirical) question to others (cf. Ludlow (2005), Stanley (2005, 69-72)).

The real problem is with premise (a). Certainly we should not posit a parameter of circumstances of evaluation without a good reason, but why suppose that the only thing that could be such a reason is an operator that

shifts the parameter? Why aren't the advantages of nonindexical contextualism over standard contextualism, recounted above, themselves a good semantic reason to posit an epistemic standards parameter?¹⁷

To see how unreasonable (a) is, consider what it would recommend if we were doing semantics for a language devoid of modal operators or counterfactual conditionals. Since this language would not contain any world-shifting operators, (a) would forbid us from relativizing propositional truth to worlds. But we would still be interested in knowing how the truth values of sentences of this impoverished language depend on features of the context of use, including the world of the context. A sentence *S* in the language—say, “Dodos were extinct in 2002”—might be true at C_1 (occurring at world w_1) and false at C_2 (occurring at world w_2). The only way we could account for this without relativizing proposition truth to worlds would be to say that *different propositions* are expressed at C_1 and C_2 . But this is highly undesirable. We would like to be able to say that a speaker at C_1 expresses the same proposition by *S* as does a speaker at C_2 , though the former speaks truly (in her context) and the latter speaks falsely (in her context). Premise (a) would forbid us from saying this, and this seems to me sufficient grounds for rejecting it, and with it Stanley's argument against an epistemic standards parameter.

Proliferation

A fourth reason for resisting an epistemic standards parameter is a worry about opening the floodgates. If “know” is context-sensitive but not indexical, it is unlikely that it is the only such expression. Very likely we'll also want nonindexical contextualist treatments of other expressions, too. To handle each new expression, we'll need a new parameter of circumstances. Pretty soon our nice ordered pairs will become ordered *n*-tuples! One might advise stopping this proliferation of parameters right at the beginning.

So stated, this isn't much of an objection. Maybe you just need a lot

¹⁷In support of his claim, Stanley says, citing Lewis (1980): “. . . the difference between elements of the circumstance of evaluation and elements of the context of use is precisely that it is elements of the former that are shiftable by sentence operators” (150). Since Lewis's “indices” are nothing more than technical devices for constructing a recursive definition of truth, it is true that there could be no motivation for positing a parameter of the index without an operator that shifts it. But, crucially, Lewis is not working with propositions in his semantic framework. (Indeed, a major point of his paper is that one does not need to.) This is an important disanalogy between his indices and Kaplan's circumstances of evaluation. The latter are constrained by factors that are irrelevant to the former, such as considerations about when we have one proposition or two distinct ones.

of parameters to do semantics. This doesn't make semantics intractable, unsystematic, or impossible (we have computers, after all). And there's no reason why we can't ignore most of these parameters when we are trying to illuminate the semantics of a particular class of expressions (say, epistemic words).

There is, however, a form of the objection that cuts more ice. By deploying "Travis cases," one might argue that even for a limited class of expressions, nonindexical contextualism would require an endless proliferation of parameters. Travis (1985) shows how context-shifting arguments can be constructed even for such apparently non-context-sensitive predicates as "weighs 160 pounds." For example, our intuitions about whether a speaker has spoken truly in saying "Smith weighs 160 pounds" will vary depending on whether Smith is about to get into a packed elevator (in which case we will count his heavy winter clothes) or at home on the bathroom scale (in which case we will want his weight naked). If we try to handle this context sensitivity the way we handled the sensitivity of "know" to epistemic standards, it seems we may end up with an about-to-board-a-crowded-elevator parameter. And now the worry is not so much that we'll have *too many* parameters, but that there will be *no end* to the addition of such parameters. The worry is that such proliferation would make systematic semantics impossible.

There are two ways the nonindexical contextualist might go here. The first would be to deny that the contextual sensitivity revealed by the Travis-style context shifting arguments is *semantic*. One might try to argue that the intuitions in these cases track speech-act content that diverges from the semantic content of the sentence type at the context of use (see Cappelen and Lepore (2005), Stainton, forthcoming). Alternatively, one might avoid the proliferation of parameters by relying on just two: a world parameter and a "counts-as" parameter, which we can model as a function from properties to *w*-intensions (functions from worlds to extensions) (see MacFarlane 2007b). The "counts-as" parameter settles what things have to be like to have various properties: e.g. the property of weighing 160 pounds, or of being tall. Because it is so all-encompassing, it eliminates the need for a proliferation of new parameters and allows us to have a systematic semantic theory.¹⁸ I won't try to decide between these options here.

¹⁸The view sketched here is similar to the one developed in Predelli (2005). Instead of countenancing an extra parameter of circumstances of evaluation, as I do here, Predelli conceives of points of evaluation as something like state descriptions (which fix the extension of every property and relation expressible in the language). Which state description is "the circumstance of the context" will depend not just on the world of the context, but on other features of context as well.

7 Utterance truth and proposition truth

A potentially more serious problem with nonindexical contextualism is the slippage it allows between proposition truth and utterance truth. Because the nonindexical contextualist holds that “know” is context-sensitive, she can accept that

(18) Janet’s utterance of “Sam knows that he has feet” in C_1 is true, and

(19) Barry’s utterance of “Sam knows that he has feet” in C_2 is not true,

where C_1 and C_2 are contemporaneous and cowardly, and Janet and Barry refer to the same person using the name “Sam.” Because she does not take “know” to be indexical, however, the nonindexical contextualist will also accept the following:

(20) At C_1 , “Sam knows that he has feet” expresses the proposition that Sam knows that he has feet (at t).

(21) At C_2 , “Sam knows that he has feet” expresses the proposition that Sam knows that he has feet (at t).

Now suppose our contextualist also accepts the following plausible principle connecting proposition truth and utterance truth:

(22) An utterance of S at C is true iff the proposition expressed by S at C is true.

Then we have trouble! For (18), (20), and (22) imply

(23) The proposition that Sam knows that he has feet (at t) is true.¹⁹

And (19), (21), and (22) imply

(24) The proposition that Sam knows that he has feet (at t) is not true.

So we have what looks like a contradiction latent in nonindexical contextualism.

The solution is obvious. The nonindexical contextualist must reject (22) in favor of

(25) An utterance of S at C is true iff the proposition expressed by S at C is true at the circumstance of C .

¹⁹The tense attached to “true” is grammatically required but semantically insignificant, since the truth value of the proposition does not vary across times of evaluation.

This follows directly from the Kaplanian definition of truth for an occurrence of a sentence in a context,

- (26) An occurrence of S in C is true iff the proposition expressed by S at C is true at the circumstance of C .

together with the bridge principle

- (27) An utterance of S at C is true iff an occurrence of S in C is true.²⁰

It is easy to explain why (22) sounds so plausible, even though it is not true in full generality. When circumstances are just worlds (as they are on the orthodox view) and the context C is located at the same world as the speaker's context, the following sentences are guaranteed to have the same truth value:

- (28) The proposition that p is true.
(29) The proposition that p is true at the circumstance of C .

For the extension of the (unrelativized) monadic predicate "true" at the speaker's context is just the set of propositions that are true at the circumstance of the speaker's context.²¹ So in these cases, (25) will imply (22).

Interestingly, even proponents of nonindexical contextualism can get mixed up here. Kompa (2002) thinks she has no choice but to accept "unpleasant sentences" like

- (30) Ascriber C says something true in uttering 'A knows that p ' but A doesn't know that p .

But if "says something true" means "expresses a proposition that is true," which I think is the most natural reading, then nonindexical contextualism does not predict that (30) is true.

Some may still feel that it is a bad consequence of nonindexical contextualism that one can consistently assert the following:

- (31) A 's utterance of S at C is true. (Or: A speaks truly in uttering S at C .)
(32) What A says in uttering S at C is false.

²⁰Kaplan's own (reasonable) view was that semanticists should not talk about utterances, since the notion of utterance belongs to the theory of speech acts. But if you do want to talk about utterance truth in semantics, (27) seems natural and unobjectionable.

²¹This follows from (26) and the following (naive, but natural) semantics for the monadic propositional truth predicate: at every context of use, the extension of "true" at a circumstance of evaluation e is the set of propositions that are true at e .

This certainly sounds weird on first hearing, but I'm not sure we should be bothered by it once we realize that utterance truth is a *technical* notion. In ordinary speech, people predicate truth of propositions (that is, of what is said or asserted or believed), not of utterances. If utterance truth is a technical notion, we had better make sure our intuitions about it are in line with our theories, not the other way around. Rejecting a theory because it makes predictions about utterance truth that "sound funny" is not sound methodology.

Indeed, it is arguable that orthodox views are already committed to the consistency of (31) and (32). Suppose that at C , A says, "Dodos were not extinct in 2002" ($=S$), and suppose that at the world of C , dodos were not extinct in 2002. Then, given our understanding of utterance truth (27), it is correct to say both

(33) A 's utterance of S at C is true, and

(34) What A says in uttering S at C [namely, that Dodos were not extinct in 2002] is false.

If you find it hard to hear (34) as true, be sure you're not confusing it with the false (and much more natural sounding) sentence

(35) What A would have said in uttering S at C *would have been* false,

in which the world of evaluation is shifted in the predicate.

For those who are still convinced that (31) and (32) are inconsistent, there is another option—namely, relativism. Whereas nonindexical contextualism lets the epistemic standard parameter be initialized by the context of use, relativism lets it be initialized by the context of assessment. On the version I favor (MacFarlane, 2005a), A 's utterance of S at C_1 is true (relative to a context of assessment C_2) just in case the proposition S expresses at C_1 is true with respect to the world of C_1 and the epistemic standards in play at C_2 . Relativism offers all of the advantages of nonindexical contextualism but does not make (31) and (32) consistent. However, it does require us to make sense of assessment-relative truth, a philosophical burden nonindexical contextualism avoids.²²

²²Whether relativism or nonindexical contextualism is correct for a discourse will depend on data about retraction and disagreement. The issues here go beyond the scope of this paper: for discussion, see MacFarlane (2005b, 2007a).

References

- Herman Cappelen and Ernie Lepore. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Blackwell, Oxford, 2005.
- Keith DeRose. Contextualism and knowledge attributions. *Philosophy and Phenomenological Research*, 52(4):913-929, 1992.
- Keith DeRose. Solving the skeptical problem. *Philosophical Review*, 104:1-52, 1995.
- Keith DeRose. Relevant alternatives and the content of knowledge attributions. *Philosophy and Phenomenological Research*, 56(1):193-197, 1996.
- Keith DeRose. Now you know it, now you don't. In *Proceedings of the Twentieth World Congress of Philosophy, Volume V: Epistemology*, pages 91-106, Bowling Green, Ohio, 2000. Philosophy Documentation Center.
- Keith DeRose. Assertion, knowledge, and context. *Philosophical Review*, 111: 167-203, 2002.
- Richard Feldman. Skeptical problems, contextualist solutions. *Philosophical Studies*, 103:61-85, 2001.
- Gottlob Frege. *Posthumous Writings*. University of Chicago Press, Chicago, 1979.
- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, Oxford, 2004.
- David Kaplan. Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In Joseph Almog, John Perry, and Howard Wettstein, editors, *Themes from Kaplan*, pages 481-566. Oxford University Press, Oxford, 1989.
- Jeffrey C. King. Tense, modality, and semantic values. *Philosophical Perspectives*, 17:195-245, 2003.
- Nikola Kompa. The context sensitivity of knowledge ascriptions. *Grazer Philosophische Studien*, 64:79-96, 2002.
- David Lewis. Index, context, and content. In Stig Kanger and Sven Öhman, editors, *Philosophy and Grammar*. Reidel, Dordrecht, 1980.

- David Lewis. *Papers in Philosophical Logic*. Cambridge University Press, Cambridge, 1998.
- Peter Ludlow. Contextualism and the new linguistic turn in epistemology. In Gerhard Preyer and Georg Peter, editors, *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, pages 11–50. Oxford University Press, Oxford, 2005.
- John MacFarlane. The assessment sensitivity of knowledge attributions. *Oxford Studies in Epistemology*, 1:197–233, 2005a.
- John MacFarlane. Making sense of relative truth. *Proceedings of the Aristotelian Society*, 105:321–339, 2005b.
- John MacFarlane. Relativism and disagreement. *Philosophical Studies*, 132:17–31, 2007a.
- John MacFarlane. Semantic minimalism and nonindexical contextualism. In G. Preyer and G. Peter, editors, *Context-Sensitivity and Semantic Minimalism*, pages 240–250. Oxford University Press, Oxford, 2007b.
- John Perry. Thought without representation. *Proceedings of the Aristotelian Society*, 1986.
- Stefano Predelli. Painted leaves, context, and semantic analysis. *Linguistics and Philosophy*, 28:351–374, 2005.
- Mark Richard. Temporalism and eternalism. *Philosophical Studies*, 39:1–13, 1980.
- Mark Richard. Introduction to part i. In Aleksandar Jokić and Quentin Smith, editors, *Time, Tense, and Reference*, pages 27–45. MIT Press, Cambridge, 2003.
- Mark Richard. Contextualism and relativism. *Philosophical Studies*, 119:215–242, 2004.
- Jonathan Schaffer. From contextualism to contrastivism. *Philosophical Studies*, 119:73–103, 2004.
- Stephen Schiffer. Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*, 96:317–333, 1996.
- Scott Soames. Replies. *Philosophy and Phenomenological Research*, 65:429–452, 2002.

- Robert Stainton. Contextualism in epistemology and the context sensitivity of 'knows'. In J. C. Campbell, M. O'Rourke, and H. Silverstein, editors, *Knowledge and Skepticism: Topics in Contemporary Philosophy*. MIT Press, Cambridge, MA, forthcoming.
- Jason Stanley. Context and logical form. *Linguistics and Philosophy*, 23:391-434, 2000.
- Jason Stanley. Context, interest relativity, and the sorites. *Analysis*, 63:269-280, 2003.
- Jason Stanley. On the linguistic basis for contextualism. *Philosophical Studies*, 119:119-146, 2004.
- Jason Stanley. *Knowledge and Practical Interests*. Oxford University Press, 2005.
- Charles Travis. On what is strictly speaking true. *Canadian Journal of Philosophy*, 15:187-229, 1985.