

Epistemic Contextualism as a Theory of Primary Speaker Meaning*

Gilbert Harman, Princeton University

June 30, 2006

Jason Stanley's *Knowledge and Practical Interests* is a brilliant book, combining insights about knowledge with a careful examination of how recent views in epistemology fit with the best of recent linguistic semantics. Although I am largely convinced by Stanley's objections to epistemic contextualism, I will try in what follows to formulate a version that might have some prospect of escaping his powerful critique.

Background

Epistemic contextualism was originally put forward as a response to a familiar skeptical Cartesian argument.

You think you know where you live, but how do you know you aren't a bodiless victim of an evil demon that has supplied you with false experiences and memories? And, if you are not in a

*I am indebted to Jason Stanley for his comments on my contribution to a symposium on his book at the 2006 Pacific Division APA meeting in Portland and to Delia Graff Fara for her comments on these remarks. I have also been greatly influenced by the conception of semantics in Neale (forthcoming).

position to know you are not such a victim, how can you know
where you live?

Dretske's (1970) *noncontextualist* answer was that the skeptic's argument wrongly assumes that knowledge claims are closed under obvious logical entailment. When we say that Albert knows where he lives, we suppose Albert *takes for granted* that he is not the victim of an evil demon, but we don't suppose he *knows* he is not such a victim. Similarly, when Albert is at the zoo, he knows that he and his child are looking at a zebra, even though he does not know that what they seeing isn't a horse painted to look like a zebra. Albert's knowledge depends on the assumption that there is no funny business of that sort going on. He knows only if that assumption is correct even if he does not know the assumption is correct.¹

Stanley (2005, p. 18), echoing many others, says about Dretske's response that its "most worrisome consequence is that it results in a straightforward failure of single-premise epistemic closure." I want to begin my own discussion by calling attention to something interesting about this reaction, namely, that the main point of Dretske's paper was precisely to argue *against* single-premise epistemic closure by appeal to intuitions about cases! Stanley's "worrisome consequence" is Dretske's intuitively plausible point in favor of his theory.

Stanley approvingly cites Keith DeRose's (1995) claim that it is particularly infelicitous to say such things as, "John knows that he has hands, but he does not know that he is not a handless brain in a vat." But such a

¹Brett Sherman and I develop a version of Dretske's response in (Harman and Sherman 2004).

remark seems quite OK to me (and to many others). When Stanley says of such examples, that Dretske's theory "has inexplicably bad consequences," it would be more appropriate to report a clash of "intuitions".

Epistemic Contextualism

In certain "ordinary" contexts people tend to allow that Albert knows where he lives, etc., while in skeptical philosophical contexts the same people are inclined not to allow such things. After those contexts have passed, however, people tend to go back to allowing what they were reluctant to allow in the skeptical philosophical contexts.

It also seems that whether speakers attribute knowledge to someone can depend on what is at stake. People seem to be more willing to allow that Hannah knows the bank will be open tomorrow if it does not matter very much to them whether the bank is open than if it does matter.

Epistemic contextualism explains some or all of these apparent facts of usage as due to a difference between what speakers mean by their knowledge attributions in ordinary low-stakes contexts and what they mean by such attributions in philosophical or high-stakes contexts. In ordinary contexts in saying that Alfred knows that he lives in New Jersey what one means in the first instance to be saying does not imply that Albert is in a position to know he is not a brain in a vat. In philosophical contexts, what one means in the first instance to be saying in saying that Alfred knows he lives in New Jersey may imply that Alfred is in such a position. In ordinary contexts, in saying Hannah knows that the bank will be open tomorrow, what one

means does not imply that her grounds meet the same standards as would be needed if a great deal were at stake.

So understood, epistemic contextualism is a claim about *speaker meaning* or *utterer's meaning* (Grice 1957, 1968). Furthermore, it is a claim about what speakers or utterers or language users *mean in the first instance to be saying* rather than what they might mean to be *implying* in saying something else. It is a claim about *primary speaker meaning*.

Importantly, epistemic contextualism understood in this way is not directly a claim about the meanings of words or sentences, although one or another theorist might go on to try to explain the supposed differences in primary speaker meaning by appeal to something special about the meaning of the word *know* or other aspects of the meaning of sentences used to make knowledge attributions.

Stanley's Critique

As an alternative to epistemic contextualism, Stanley offers the plausible principle that whether *S* knows that *P* can depend on the importance to *S* not to have a false belief that *P*. Whether *S* knows that *P* can depend on whether *S* is justified in deciding that the grounds for accepting *P* are sufficient, which can depend on how important it is to *S* not to have a false belief that *P*.

While I agree with Stanley's principle (Harman 1986 and 2004), it does not by itself adequately account for all the data. For one thing, being in skeptical philosophical contexts does not normally seem to have the relevant

sort of effect on practical interests. Furthermore, Stanley's principle does not by itself explain why the truth or appropriateness of a judgment about what Albert knows would be affected by whether or not the *person making that judgment* (rather than Albert) is or is not in a philosophical context or by how important it is that the person making the judgment not have a false belief in this instance.

Stanley agrees he needs some additional assumptions to account for the data and he goes on to argue that his resulting view with those assumptions is to be preferred to any sort of epistemic contextualism. In particular, he argues impressively that no existing contextualist explanation of the data is plausible in the light of contemporary linguistic semantics. He demonstrates in detail the failure of various alleged parallels between epistemic contextualism and other more familiar sort of contextualism involving references of pronouns, quantifier domain restrictions, comparison classes for adjectives like *rich*, and many others.

Minimal Contextualism

Nevertheless, I think it might be possible to defend a *minimalist contextualism* that starts with the idea that what a person knows typically rests on things the person simply assumes or takes for granted without knowing them and without being in a position to know them. For example, a person typically takes for granted that the world is much as it seems and is not an illusion created by a Cartesian evil demon. A visitor to the zoo typically assumes that the animals displayed there have not been painted or otherwise

disguised so as to look like other animals.

In this view, when others think about what people know, they too typically take for granted the same sorts of things. When they do so, they are perfectly willing to say such things as that Albert knows he lives in New Jersey and knows that the animal he is looking at is a zebra.

Minimal contextualism holds that in such a case neither what someone means in the first instance to say in saying that Albert knows he lives in New Jersey nor what the audience takes him or her to mean in the first instance to be saying implies that Albert knows he is not a brain in a vat or other evil demon victim. Similarly, what someone means in the first instance to say in saying that Albert knows he is looking at a zebra does not imply, and would not be taken to imply, that he knows that it isn't a horse painted to look like a zebra. Those are things that a speaker and other participants in the conversation typically take Albert to assume without necessarily knowing to be the case, although conversational participants should not be expected to be able to express matters in exactly this way.

Furthermore, what the speaker means in the first instance to be saying in these contexts can be true even though Albert merely assumes and does not know that the relevant presuppositions are true.

On the other hand, there are other contexts in which discussants do not share Albert's presuppositions, contexts in which what someone means in the first instance to be saying in saying that Albert knows he lives in New Jersey *does* imply Albert knows he is not a mere victim of an evil demon, and contexts in which what someone means in the first instance to be saying in saying that Albert knows he sees a zebra *does* imply Albert knows it is

not a horse painted to look like a zebra. Furthermore, what a speaker means in the first instance to be saying in these contexts can be false, if Albert does not know he is not a mere victim of an evil demon or does not know he isn't seeing a horse painted to look like a zebra.

My first question, then, is whether this minimal epistemic contextualism about primary speaker meaning can already be shown to be in trouble. If not, I have some thoughts about how to develop a slightly less minimal version.

Slightly Less Minimal Contextualism

If the attributer is in a *skeptical philosophical* context but Albert is in an ordinary context, then when the attributer says that Albert does not know he lives in New Jersey because Albert does not know he isn't a victim of an evil Cartesian demon, what the attributer means in the first instance to be saying might be true.

What if instead the attributer is in a *high practical stakes* context but the subject is not? The attributer says Hannah does not know the bank will be open tomorrow. It would be disastrous for the attributer (but not for Hannah) to believe the bank will be open tomorrow if it won't be. Does the truth value of what the attributer means in the first instance to be saying depend on whether Hannah's belief meets the high standards that would be required by the attributer's high stakes? I am inclined to think so, but I am not sure.

What if Albert is in a skeptical philosophical context but the attributer,

Mabel, is not? Mabel says Albert knows he lives in New Jersey. Albert says he does not know he lives in New Jersey because he does not know that he isn't a brain in a vat. I am not sure about this but it seems to me that what Mabel means in the first instance to be saying can be true in such a case. On the other hand, what Albert means in the first instance to be saying in his context when he says that he does not know he lives in New Jersey can also be true, it seems to me. What Albert means in the first instance in his skeptical context does not it seems to me have to conflict with what is meant in the first instance by Mabel in her ordinary nonphilosophical context.

What if instead Hannah is in a high practical stakes context but Mabel is not? Mabel says Hannah knows the bank will be open tomorrow. It happens to be very important for Hannah but not for Mabel not to believe mistakenly that the bank will be open if it won't be. Here I agree with Stanley that whether what Mabel means in the first instance to be saying is correct depends on whether Hannah knows in relation to the high stakes, because whether Hannah's justification for believing the bank will be open is sufficient depends on what is at stake.

So, I am inclined to see a difference between way in which subject high stakes are relevant and the way in which skeptical subject contexts are relevant to knowledge attributions.

Standards of Precision?

Lewis (1979) suggests that (what I am calling) minimal epistemic contextualism might be seen as an instance of variation in presupposed "standards

of precision.” For example, someone says that there is a vacuum in a certain container (e.g., in a vacuum tube). The truth of what the speaker means in the first instance to be saying depends on what standards the speaker presupposes. Given the strictest imaginable standards, perhaps there are no vacuums, because nothing is completely empty, using the strictest standards of emptiness. Given laxer standards, there is a vacuum in this particular tube.

Similarly perhaps, by lax standards, Albert knows he lives in New Jersey; by stricter standards he does not know this. What a speaker means in the first instance to be saying in saying that Albert knows he lives in New Jersey is true if the speaker is presupposing lax standards but not true if the speaker is presupposing stricter standards.

Objection: this is the sort of “loose use” involved in saying that France is hexagonal. It is not a literal use of language. It is intended to be recognized as an exaggeration. The speaker literally says something false but implies something more or less true. To take that as a model for epistemic contextualism is to accept Unger’s (1975) skepticism. In that view, no one ever really knows anything, but we communicate something indirectly by saying people know things just as we communicate something about France in saying that it is hexagonal even though France is not really hexagonal.

Reply: Perhaps that is right about France not really being hexagonal, but without further argument the objection begs the question against this version of epistemic contextualism, which asserts that knowledge claims are not, and are not intended to be, understood as exaggerations. Epistemic contextualism is here understood as a claim about direct speaker meaning,

not a theory of what the speaker intends to be implying or implicating or otherwise saying indirectly.

Bibliography

DeRose, Keith, (1995). "Solving the Skeptical Puzzle," *Philosophical Review* 104: 1-52.

Dretske, Fred, (1970). "Epistemic Operators," *Journal of Philosophy* 67, pp. 1007-1023.

Grice, H. P., (1957). "Meaning," *Philosophical Review* 64: 377-388.

Grice, H. P., (1968). "Utterer's Meaning, Sentence-Meaning and Word-Meaning," *Foundations of Language* 4: 225-42.

Harman, Gilbert, (1986). *Change in View: Principles of Reasoning* (Cambridge, Massachusetts: M.I.T. Press).

Harman, Gilbert, (2004). "Practical Aspects of Theoretical Rationality," in *The Oxford Handbook of Rationality*, Al Mele and Piers Rawling, eds., Oxford: Oxford University Press, pp. 45-56.

Harman, Gilbert, and Sherman, Brett, (2004). "Knowledge, assumptions, lotteries," *Philosophical Issues*, 14, pp. 492-500.

Lewis, David, (1979). "Scorekeeping in a Language Game," *Journal of Philosophical Logic* 8, pp. 339-359.

Neale, Stephen, (forthcoming). “On Location.” To appear in M. O’Rourke and C. Washington (eds.), *Situating Semantics: Essays in Honour of John Perry*. Cambridge: MIT Press, 2006.

Stanley, Jason, (2005). *Knowledge and Practical Interests* (Oxford: Clarendon Press).

Unger, Peter, (1975). *Ignorance*, Oxford: Oxford University Press.